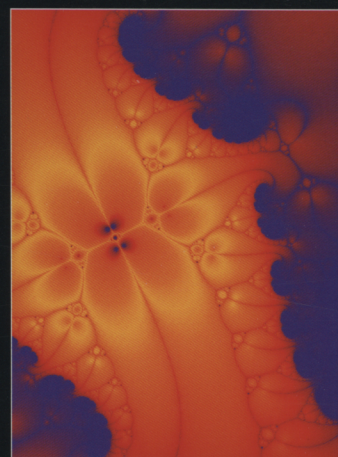




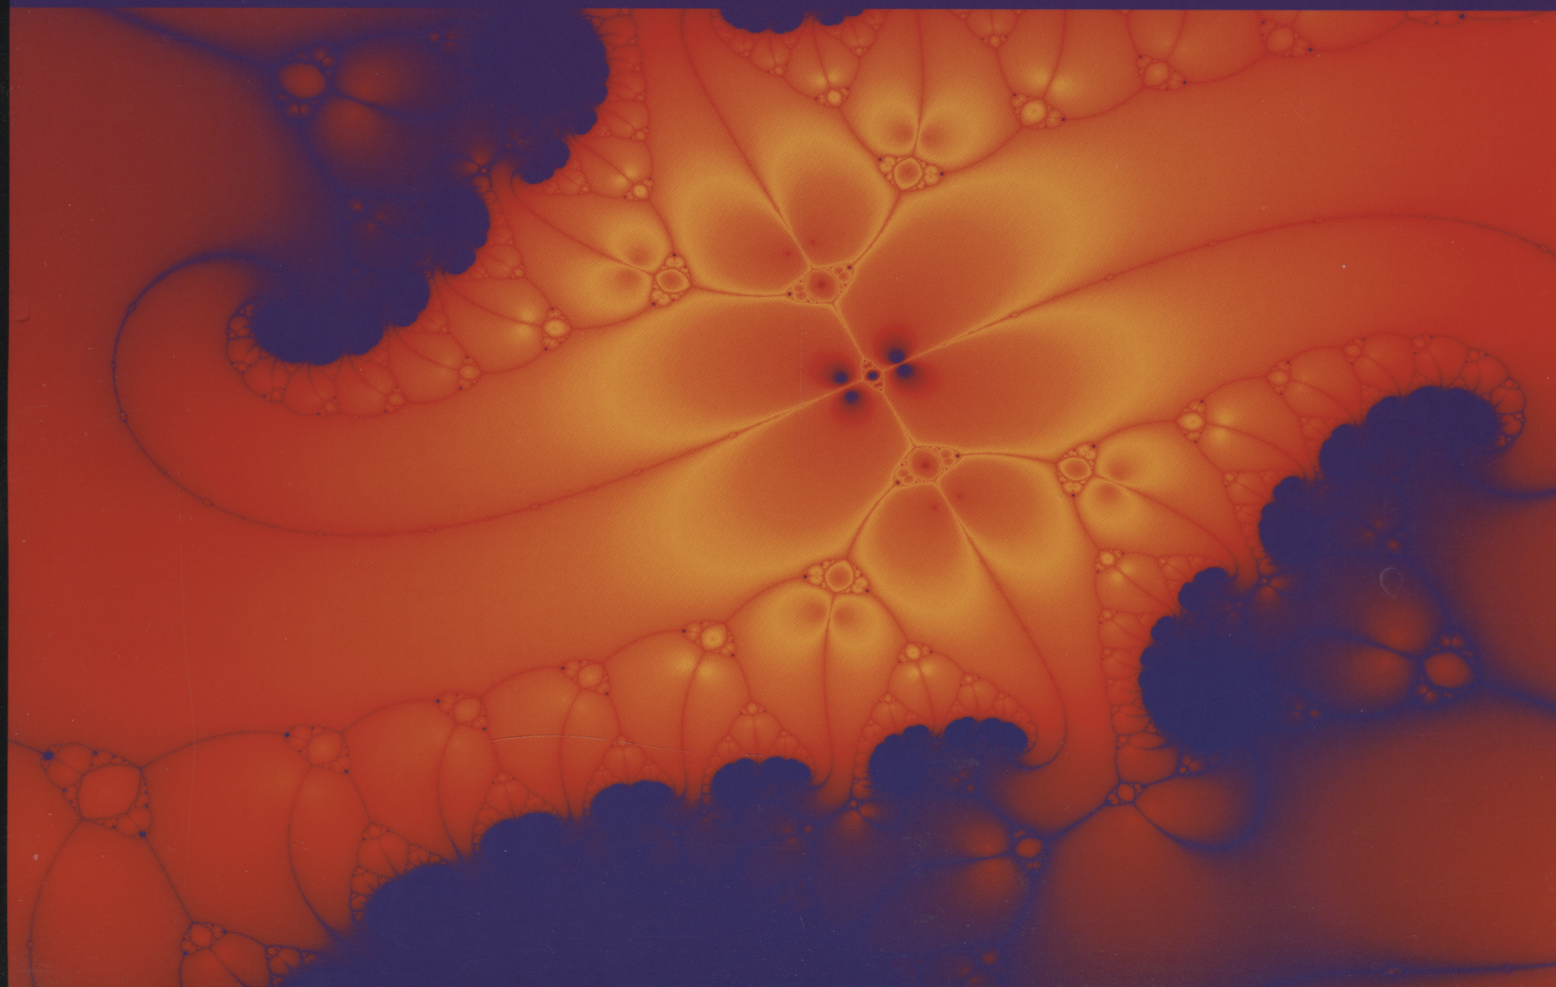
The Open University

M338 Topology

B1



Unit B1 Surfaces





The Open University

M338 Topology

B1

Surfaces

This publication forms part of an Open University course. Details of this and other Open University courses can be obtained from the Student Registration and Enquiry Service, The Open University, PO Box 197, Milton Keynes, MK7 6BJ, United Kingdom: tel. +44 (0)870 333 4340, e-mail general-enquiries@open.ac.uk

Alternatively, you may visit the Open University website at <http://www.open.ac.uk> where you can learn more about the wide range of courses and packs offered at all levels by The Open University.

To purchase a selection of Open University course materials, visit the webshop at www.ouw.co.uk, or contact Open University Worldwide, Michael Young Building, Walton Hall, Milton Keynes, MK7 6AA, United Kingdom, for a brochure: tel. +44 (0)1908 858785, fax +44 (0)1908 858787, e-mail ouwenq@open.ac.uk

The Open University, Walton Hall, Milton Keynes, MK7 6AA.

First published 2006.

Copyright © 2006 The Open University

All rights reserved; no part of this publication may be reproduced, stored in a retrieval system, transmitted or utilised in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without written permission from the publisher or a licence from the Copyright Licensing Agency Ltd. Details of such licences (for reprographic reproduction) may be obtained from the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP.

Open University course materials may also be made available in electronic formats for use by students of the University. All rights, including copyright and related rights and database rights, in electronic course materials and their contents are owned by or licensed to The Open University, or otherwise used by The Open University as permitted by applicable law.

In using electronic course materials and their contents you agree that your use will be solely for the purposes of following an Open University course of study or otherwise as licensed by The Open University or its assigns.

Except as permitted above you undertake not to copy, store in any medium (including electronic storage or use in a website), distribute, transmit or re-transmit, broadcast, modify or show in public such electronic materials in whole or in part without the prior written consent of The Open University or in accordance with the Copyright, Designs and Patents Act 1988.

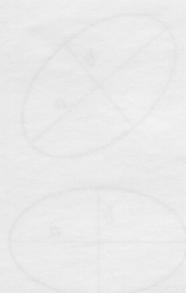
Edited, designed and typeset by The Open University, using the Open University TeX System.

Printed and bound in the United Kingdom by The Charlesworth Group, Wakefield.

ISBN 0 7492 4129 2

Contents

Introduction	4
Study guide	6
1 Examples of surfaces	7
1.1 Surfaces in space	7
1.2 Examples of surfaces	11
1.3 Paper-and-glue constructions	15
1.4 Homeomorphic surfaces	19
1.5 Defining surfaces	22
2 The orientability of surfaces	24
2.1 Surfaces with twists	24
2.2 Orientability	27
2.3 The projective plane	30
3 The Euler characteristic	35
3.1 Nets on surfaces	35
3.2 Subdivisions	37
3.3 The Euler characteristic	40
3.4 Some general results	44
3.5 The Classification Theorem	46
4 Edge identifications	50
4.1 The identification topology	50
4.2 Neighbourhoods	53
Solutions to problems	56
Index	60



Introduction

This block is concerned with a special class of topological spaces called *surfaces*. Common examples of surfaces are the sphere and the cylinder; less common, though probably still familiar, are the torus and the Möbius band. Other surfaces, such as the projective plane and the Klein bottle, may be unfamiliar, but they crop up in many places in mathematics. Our aim is to *classify* surfaces — that is, to produce criteria that allow us to determine whether two given surfaces are homeomorphic.

Recall from *Unit A3* that two topological spaces (X, \mathcal{T}_X) and (Y, \mathcal{T}_Y) are *homeomorphic* if there is a bijection $f: X \rightarrow Y$ that is continuous, and whose inverse f^{-1} is also continuous, with respect to the given topologies; such a function f is called a *homeomorphism*. The relation ‘is homeomorphic to’ between topological spaces is the most fundamental relation in topology, because two topological spaces that are homeomorphic are indistinguishable from a topological point of view — in the terminology of *Unit A3*, they are topologically equivalent. This can be seen from the alternative definition of a homeomorphism given in *Unit A3*: a bijection $f: X \rightarrow Y$ is a homeomorphism if and only if both f and f^{-1} map open sets to open sets. Thus, if (X, \mathcal{T}_X) and (Y, \mathcal{T}_Y) are homeomorphic, then not only are the elements of X and Y in one-one correspondence, but so also are their open sets. We can thus regard (Y, \mathcal{T}_Y) as being essentially the same space as (X, \mathcal{T}_X) , so far as its purely topological properties are concerned: (X, \mathcal{T}_X) and (Y, \mathcal{T}_Y) are merely two different ways of presenting the same space.

It is therefore useful to be able to determine whether two given topological spaces are homeomorphic. Of course, if we can find a specific homeomorphism between them then the question is answered; but if we fail to find a homeomorphism then we cannot deduce that there isn’t one. So looking for a particular homeomorphism may not be the best approach. What we want are criteria that enable us to tell whether two given surfaces are homeomorphic without our having to try to construct a specific homeomorphism between them.

A similar situation, which may help to clarify what is involved, is the following. A map between metric spaces that preserves distances is called an *isometry*, and two subsets of a metric space that are related by an isometry are essentially the same, so far as their metrical properties are concerned. For example, consider the collection of ellipses in \mathbb{R}^2 , with the Euclidean metric. Two ellipses are isometric if and only if the lengths of their major axes are the same, and the lengths of their minor axes are the same (see Figure 0.1). These two criteria enable us to determine whether two given ellipses are isometric without our having to construct a specific isometry between them: we simply compare the lengths of their axes. We say that the lengths of the major and minor axes *classify ellipses in \mathbb{R}^2 up to isometry*.

Notice that, in the phrase ‘classify ellipses in \mathbb{R}^2 up to isometry’, we have specified the class of figures under consideration (ellipses in \mathbb{R}^2): our criteria apply only to these particular figures in the plane, and we make no claim that other figures can be classified up to isometry in a similar way. Note also the importance of including the qualification ‘up to isometry’, since *all* ellipses are homeomorphic.

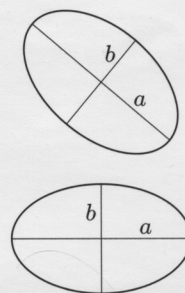


Figure 0.1 Isometric ellipses

In this block, our main task is to define a certain class of surfaces called *compact surfaces*, and then to specify criteria that allow us to determine whether two given compact surfaces are homeomorphic. These surfaces can be classified using just three criteria, as we shall show. Indeed, just as we can classify ellipses in \mathbb{R}^2 up to isometry by specifying just two numbers (the lengths of the major and minor axes), so we can classify all compact surfaces up to homeomorphism by specifying just three numbers. The appropriate theorem — the *Classification Theorem* — is stated in this unit and proved in *Unit B3*. In this unit, we begin the build-up to the proof of the Classification Theorem by giving examples of surfaces, by showing how they can be manipulated, defined and represented, and by examining some of their important properties.

The spirit in which we explore surfaces differs somewhat from the way we approached topology in Block A. This arises from the way that the theory of surfaces relates to the rest of topology, and that in turn is reflected in the history of these subjects. When studying surfaces we can proceed a long way with intuitive concepts, and the subject is often studied in that fashion. Today we may appeal to these intuitions, secure in the knowledge that the intuitively appealing theorems can be proved rigorously. Indeed, the precise definitions and rigorous proofs appearing in the literature frequently involve subtle and difficult arguments of a kind otherwise out of keeping with geometric topology, and may add little to our understanding of the underlying concepts.

An example is the so-called *Jordan Curve Theorem*, which asserts that any curve in the plane that is homeomorphic to a circle divides the plane into two regions, one bounded (the ‘inside’ of the curve) and the other unbounded (the ‘outside’); Figure 0.2 shows such a curve C in the plane, with its inside shaded. Although the *statement* of the Jordan Curve Theorem is entirely plausible — indeed, it is fundamental to many parts of mathematics — the theorem is difficult to prove. In fact, it is likely that few mathematicians have ever seen a proof, because such proofs merely confirm what we already believe. It is good to have a proof, of course, but we do not consider it essential to prove all of the results in this course. Examining the proof of the Jordan Curve Theorem reveals no particularly counter-intuitive complications, and so we omit the proof.

In this block we indicate when concepts are to be understood at an intuitive level, and when more precise definitions, theorems and proofs are required. Where possible, we provide rigorous definitions and proofs.

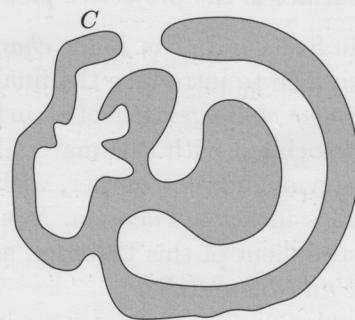


Figure 0.2

Study guide

The topology of surfaces is a highly visual subject, and this block is accompanied by a supply of examples on the DVD and in the software. You should watch the DVD again before studying this unit, as it illustrates several ideas that you will meet here. You should also find the animations in the software for this unit a useful resource for obtaining a deeper understanding of the surfaces discussed.



In Section 1, *Examples of surfaces*, we describe surfaces informally and discuss various methods for constructing them, such as the identification of edges that you saw in the DVD. We also introduce the *boundary number* of a surface — one of the three numbers in the Classification Theorem. It is important that you understand the ideas in this section before proceeding to later sections, and for this purpose you may find the software animations useful.

In Section 2, *Non-orientable surfaces*, we continue our discussion of surfaces by investigating surfaces with a ‘twist’ in them, such as the Möbius band. This leads to another number in the Classification Theorem — the *orientability number* of a surface. An important non-orientable surface is the *projective plane*, and we study its properties in some detail.

In Section 3, *The Euler characteristic*, we use the idea of *subdividing* a surface to introduce the final number in the Classification Theorem — the *Euler characteristic* of a surface. There is a short software activity associated with this material. We conclude the section by pulling the previous ideas together, with a statement and discussion of the Classification Theorem. You should make sure that you understand the statement of this theorem, as it will form the basis of our investigations in *Units B2* and *B3*.



Finally, in Section 4, *Edge identifications*, we discuss in greater detail the mathematical ideas underlying the construction of compact surfaces. This section is more technical than the previous ones, and you may prefer simply to skim through it if you are short of time. This section is not assessed.

1 Examples of surfaces

After working through this section, you should be able to:

- ▶ explain the terms *surface*, *surface in space*, *disc-like neighbourhood* and *half-disc-like neighbourhood*;
- ▶ explain the terms *n-fold torus*, *torus with n holes*, *Möbius band* and *Klein bottle*;
- ▶ explain what is meant by the *boundary* of a surface, and determine the *boundary number* of a given surface with boundary;
- ▶ construct certain *compact surfaces* from a polygon by identifying edges;
- ▶ explain how a surface in space can be regarded as a topological space.

In this section we start by introducing surfaces informally, considering several familiar examples such as the sphere, cube and Möbius band. We also illustrate how surfaces can be constructed from a polygon by identifying edges. A more formal approach to surfaces is presented at the end of the section.

1.1 Surfaces in space

Figure 1.1 shows some simple examples of surfaces in three-dimensional space.

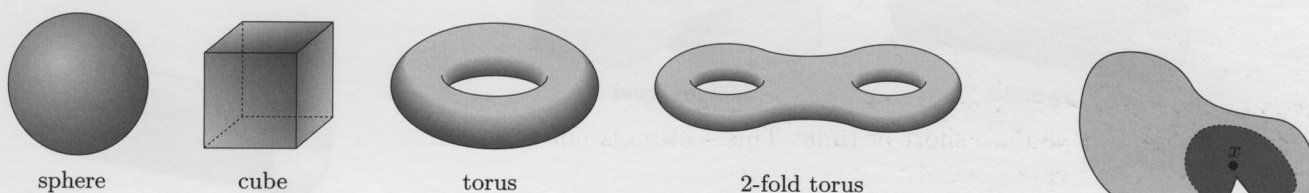


Figure 1.1

Since we are interested only in the surface of such objects, rather than the interior, you should always think of such objects as being hollow, rather than solid. Such surfaces, being subsets of \mathbb{R}^3 , are given the subspace topology from \mathbb{R}^3 .

In each of the examples in Figure 1.1, the surface is *locally* like part of the plane, as illustrated in Figure 1.2. This means that, for each point x of the surface, we can find a subset of the surface that contains x and is homeomorphic to an open disc in the plane: we call such a subset a **disc-like neighbourhood** of x . Such a neighbourhood is an open subset of the surface in the subspace topology that it inherits from \mathbb{R}^3 .

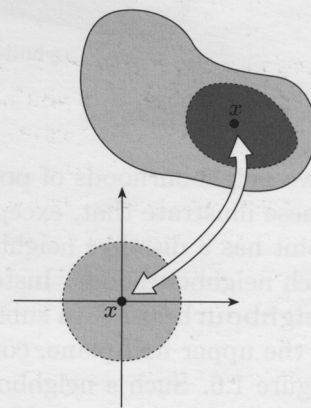


Figure 1.2 Disc-like neighbourhood of x

The subspace topology was defined in Unit A3.

Recall, from Unit A4, that a *neighbourhood* of a point x is an open set containing x .

Figure 1.3 shows some disc-like neighbourhoods in three of the surfaces of Figure 1.1. Notice that a disc-like neighbourhood may be curved or bent in space, or may even have folds in it.

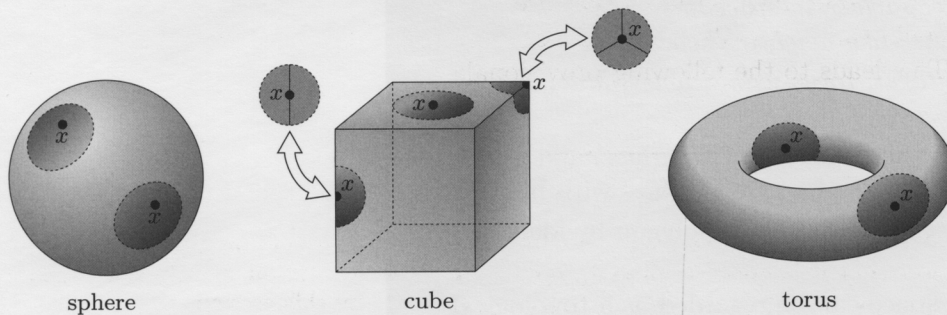


Figure 1.3

Problem 1.1

Draw some disc-like neighbourhoods on a 2-fold torus.

All the objects we wish to consider as surfaces have many points with disc-like neighbourhoods. However, there are some surfaces where some of the points do not have such neighbourhoods. Two common examples are shown in Figure 1.4 — a hollow **cylinder**, open at both ends, and a **Möbius band**.

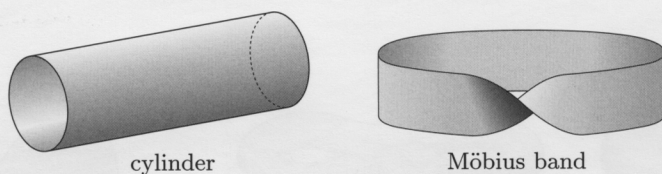


Figure 1.4

Two neighbourhoods of points on the cylinder are shown in Figure 1.5. These illustrate that, except for the points at the ends of the cylinder, each point has a disc-like neighbourhood; but the points at the ends have no such neighbourhoods. Instead, each such endpoint x has a **half-disc-like neighbourhood** — a subset that is homeomorphic to an **open half-disc** in the upper half-plane, considered as a subspace of \mathbb{R}^2 , as shown in Figure 1.6. Such a neighbourhood is again an open subset of the surface in the subset topology that it inherits from \mathbb{R}^3 .

Figure 1.5

The upper half-plane includes the points on the horizontal axis. The open half-discs are those with diameter along that axis and include the axis points; they are open subsets of the closed half-plane in the subspace topology inherited from \mathbb{R}^2 .

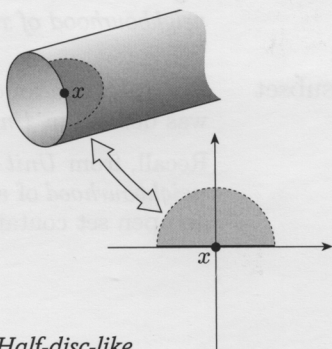


Figure 1.6 Half-disc-like neighbourhood of x

Problem 1.2

Draw some disc-like and half-disc-like neighbourhoods on a Möbius band.

In fact, all the objects we wish to consider have surfaces that are characterized as consisting of points that have either disc-like or half-disc-like neighbourhoods. This leads to the following provisional definition.

Provisional definition

A **surface** is a topological space (X, \mathcal{T}) with the property that each point of X has a neighbourhood that is homeomorphic either to an open disc in the plane or to an open half-disc in the upper half-plane, i.e. given any point $x \in X$, there is an open set U containing x such that U is homeomorphic either to an open disc in \mathbb{R}^2 or to an open half-disc in the upper half-plane with the subspace topology inherited from the Euclidean topology on \mathbb{R}^2 . The open set U is called a **disc-like neighbourhood**, or **half-disc-like neighbourhood** of x , as the case may be.

We shall modify this definition slightly at the end of this section.

This idea of points of a surface having disc-like neighbourhoods or half-disc-like neighbourhoods is crucial to the concept of a surface. In the following worked problem we present an object that looks as though it might be a surface, but fails to be so because it has a point with no such neighbourhood.

Worked problem 1.1

Explain why the double cone, open at both ends, is not a surface.

Solution

Although most points on the double cone have disc-like or half-disc-like neighbourhoods, the point z in the middle, indicated in Figure 1.8(a), does not. One way to see this is to note that if we remove the point z , then the neighbourhood breaks into two pieces, as shown in Figure 1.8(b). This cannot be the case if the neighbourhood is homeomorphic to a disc or half-disc: removing a single point from a disc or half-disc cannot split it into two pieces.

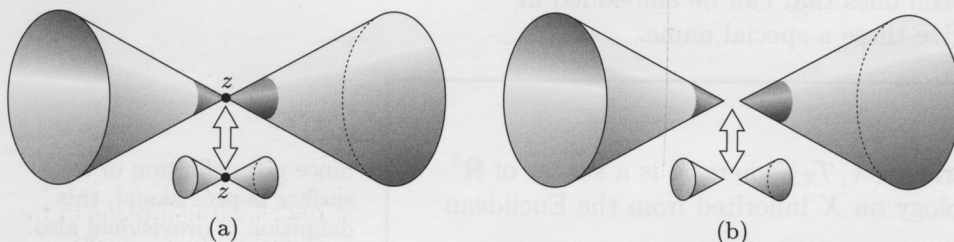


Figure 1.8

Since we have found a point with no disc-like or half-disc-like neighbourhood, the double cone is not a surface. ■

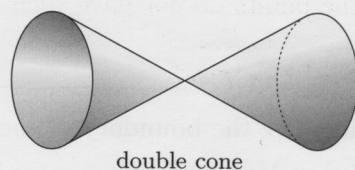


Figure 1.7

Problem 1.3

Explain why the 3-leaved book, illustrated in Figure 1.9, is not a surface.

The points of a surface with half-disc-like neighbourhoods, such as the endpoints of a cylinder, can be thought of as lying on the *boundary* of the surface.

Definitions

Let $S = (X, \mathcal{T})$ be a surface.

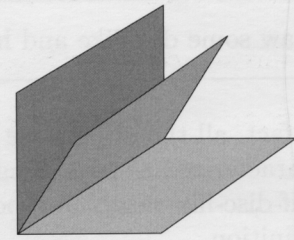
A **boundary point** of S is a point of X for which every neighbourhood contains a half-disc-like neighbourhood.

The **boundary** of S is the set of all boundary points of S .

A **surface with boundary** is a surface whose boundary is non-empty.

A **surface without boundary** is a surface with no boundary points.

An example of a surface with boundary is a cylinder, whose boundary consists of the two circular ends (Figure 1.10). An example of a surface without boundary is a sphere, since each point has a disc-like neighbourhood (Figure 1.3).



3-leaved book

Figure 1.9

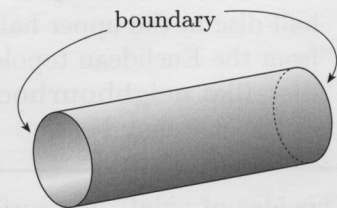


Figure 1.10

Problem 1.4

Identify the boundary of each of the following surfaces:

- (a) a Möbius band;
- (b) a torus.

The following theorem describes the boundary of a surface.

Theorem 1.1

The boundary of a surface is composed of a finite number of arcs, each of which is homeomorphic to a circle or to a unit interval, or it is the empty set.

We omit the proof, which is beyond the scope of this course.

Up to now our surfaces have been ones that can be embedded in three-dimensional space. We give these a special name.

Provisional definition

A **surface in space** is a surface (X, \mathcal{T}_X) where X is a subset of \mathbb{R}^3 and \mathcal{T}_X is the subspace topology on X inherited from the Euclidean topology \mathcal{T} on \mathbb{R}^3 .

Since our definition of a surface is provisional, this definition is provisional also.

In Subsection 1.3 we shall meet some surfaces that cannot be embedded in three-dimensional space.

1.2 Examples of surfaces

In this subsection we present a wide range of examples of surfaces in space.

Surfaces without boundary

Examples of surfaces without boundary are a sphere and a torus. Other examples are the following.

n -fold toruses

Figure 1.11 depicts a 2-fold torus and a 3-fold torus, with two and three rings respectively. An n -fold torus, for any positive integer n has n rings.

A 1-fold torus is simply a torus, as in Figure 1.1.

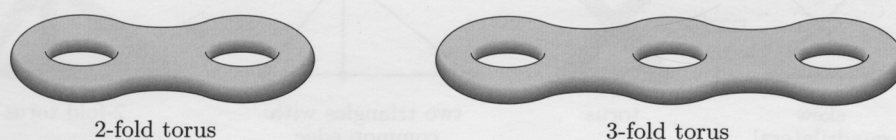


Figure 1.11

Toruses do not have to be considered as ' n -ringed doughnuts' as in Figure 1.11; they can also usefully be described as 'spheres with n handles', as Figure 1.12 illustrates.

In fact, any homeomorphic representation will do, as we shall see in Subsection 1.4.

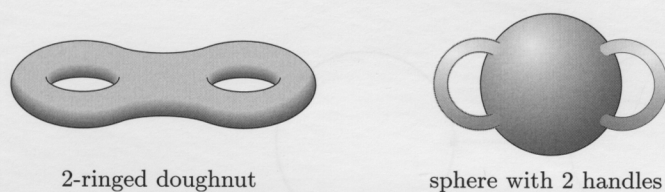


Figure 1.12 Two different representations of a 2-fold torus

Hollow tubing surfaces

In their doughnut-shaped representation, toruses can be thought of as being hollow tubes. Many other surfaces in space can also be drawn as if they were made of hollow tubing. Figure 1.13 shows two such examples.

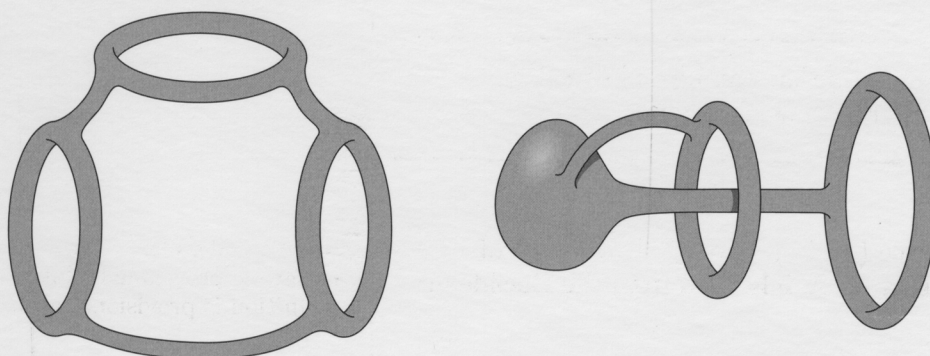


Figure 1.13

There is an easy way of constructing hollow tubing surfaces. Start with a finite set of points in three-dimensional space, called **vertices**, and connect them together by a finite set of lines, called **edges**, that meet only at the vertices, as shown in Figure 1.14: such a configuration is called a **connected graph**. We then obtain a surface by 'thickening each edge' of the connected graph into a hollow tube, as illustrated in Figure 1.15: the resulting surface, called a **thickened graph**, can be thought of as a hollow tube wrapping the original connected graph.

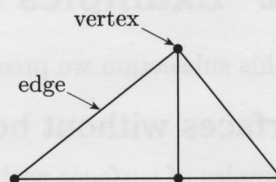


Figure 1.14 Connected graph with 4 vertices and 5 edges

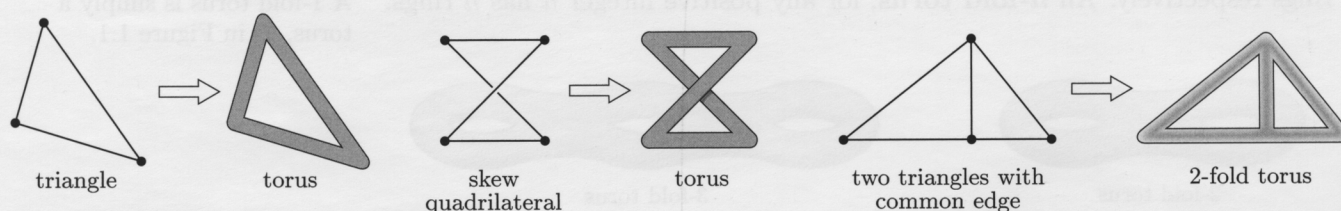


Figure 1.15

Problem 1.5

For each of the connected graphs in Figure 1.16, draw the associated thickened graph.

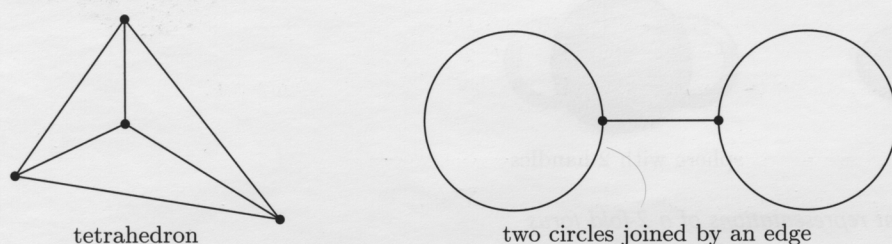


Figure 1.16

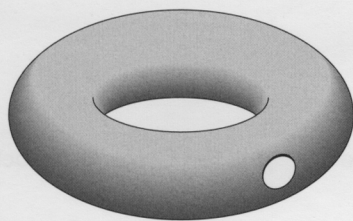
Surfaces with boundary

Examples of surfaces with boundary are a cylinder and a Möbius band. Other examples are the following.

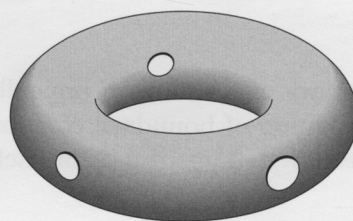
Surfaces with holes

We can obtain a surface with boundary by taking any surface without boundary and punching some holes in it by removing open discs. For example, Figure 1.17 shows a sphere with 3 holes. The boundary of the surface with holes consists of the boundaries of the holes.

Another example is a **torus with holes**. Figure 1.18 illustrates a torus with 1 hole and a torus with 3 holes. We can have a torus with k holes for any positive integer k .



torus with 1 hole



torus with 3 holes

Figure 1.18

We can also combine the notions of an n -fold torus and a torus with k holes to draw an n -fold torus with k holes, for any positive integers n and k . For example, Figure 1.19 depicts a 2-fold torus with 4 holes.

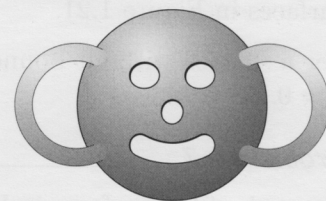


Figure 1.19

Problem 1.6

Draw a 4-fold torus with 5 holes and a 5-fold torus with 4 holes. What is the boundary of each surface?

Paper surfaces

Other surfaces can be drawn as if they were made from thin strips of material, such as paper. The surface consists of one side of the paper only. Two such surfaces are shown in Figure 1.20. Notice how some parts of a surface can cross over, or under, other parts.

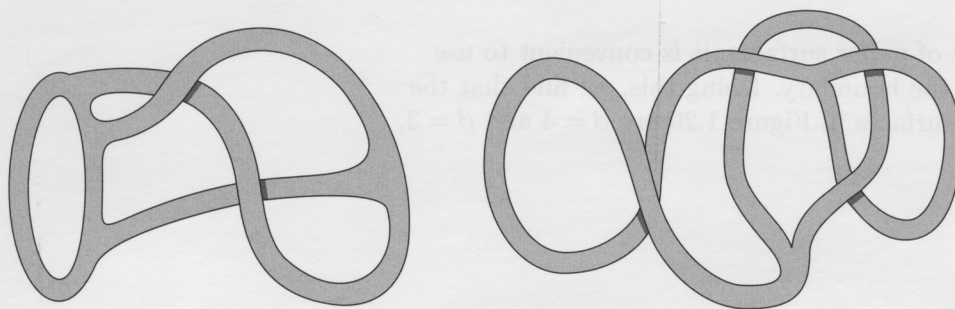


Figure 1.20 Two paper surfaces

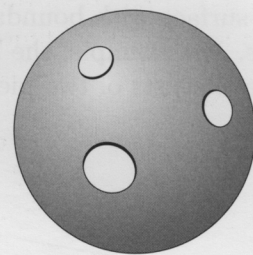


Figure 1.17

Do not confuse a torus with k holes and a k -fold torus.

Boundary numbers

In a surface with boundary, the boundary may consist of more than one piece. For example, the boundary of a cylinder, or of a torus with two holes, consists of two pieces, as shown in Figure 1.21.

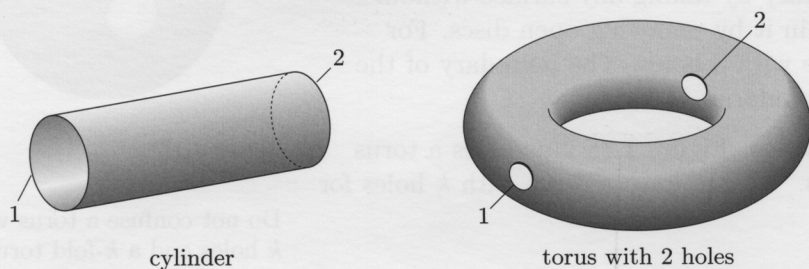


Figure 1.21

The separate pieces forming the boundary of a surface S are called the **boundary components** of S . The number of boundary components is the **boundary number** of S , denoted by β . Thus $\beta = 2$ for each of the surfaces in Figure 1.21.

β is the Greek letter beta.

For a surface without boundary, there are no boundary components, so $\beta = 0$.

Problem 1.7

For each of the surfaces in Figure 1.22, write down the boundary number β .

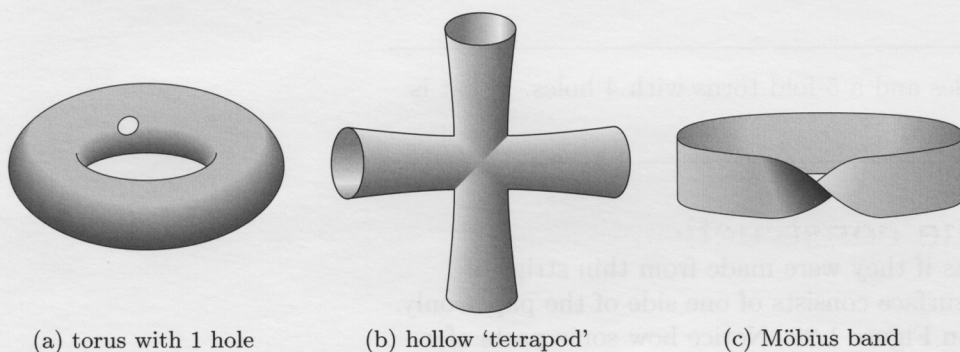


Figure 1.22

To find the boundary numbers of paper surfaces, it is convenient to use coloured pens to trace around the boundary. Doing this, we find that the two boundary numbers of the surfaces in Figure 1.20 are $\beta = 4$ and $\beta = 3$, as shown in Figure 1.23.

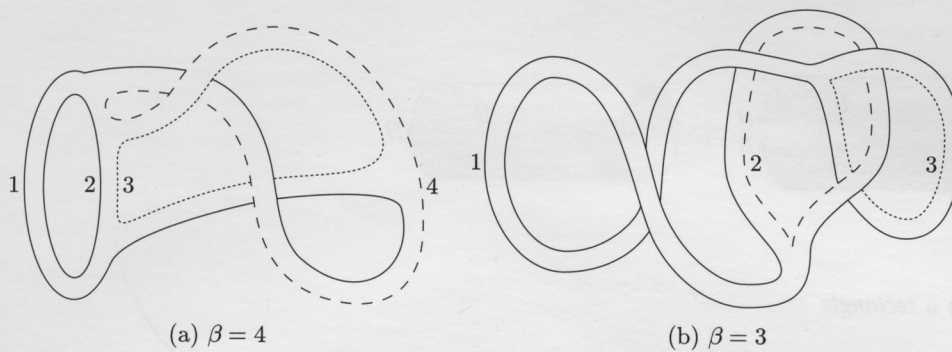


Figure 1.23

Problem 1.8

Find the boundary number β of each of the paper surfaces in Figure 1.24.

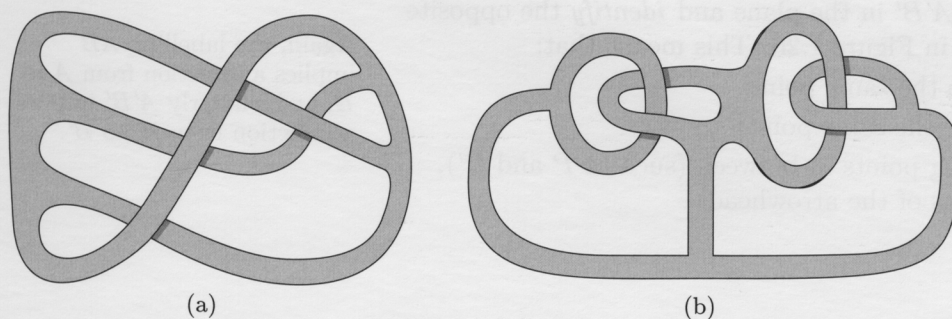


Figure 1.24

The boundary number β is one of the three characteristic numbers that we use to classify compact surfaces.

1.3 Paper-and-glue constructions

In this subsection we show how to construct surfaces by taking a piece of paper in the shape of a polygon and gluing some of its edges together. You have already seen examples of this on the DVD, and several of the surfaces in this subsection appear on the software. The surfaces that we obtain occupy a central position in this block, as you will see.

Cylinder

The simplest example of a paper-and-glue construction is to make a rectangle into a *cylinder* by gluing together two opposite edges. We take a closed rectangle $ABB'A'$ in the plane and *identify* the opposite edges AB and $A'B'$, as shown in Figure 1.25. This means that:

- we imagine A and A' to be the same point (which we'll label A);
- we imagine B and B' to be the same point (which we'll label B);
- we match up all corresponding pairs of points in between (such as P and P') as shown.

Note that we have drawn arrowheads on the edges AB and $A'B'$ of the rectangle to emphasize that they are to be identified *in the same direction*.



The direction of the identification is given by the order of the letters labelling an edge: labelling an edge AB implies a direction from A to B ; similarly $A'B'$ implies a direction from A' to B' .

We also obtain a cylinder if we identify the edges AA' and BB' .

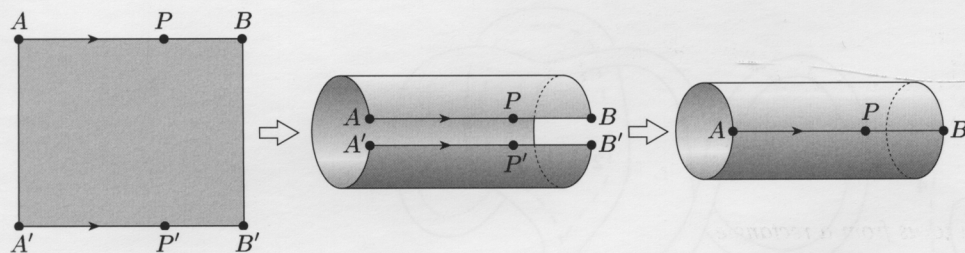


Figure 1.25 Making a cylinder from a rectangle

Möbius band

Let us now see what happens if we try to identify two opposite edges of a rectangle in *opposing* directions. We start with a rectangle, as before, but this time the edges to be identified have their arrowheads pointing in opposite directions. This means that we cannot glue the edges directly, but have to twist one of them through π radians before gluing. More formally, we take a closed rectangle $ABA'B'$ in the plane and *identify* the opposite edges AB and $A'B'$, as shown in Figure 1.26. This means that:

- we imagine A and A' to be the same point;
- we imagine B and B' to be the same point;
- we pair up all corresponding points in between (such as P and P'), taking note of the directions of the arrowheads.

We obtain a *Möbius band*.

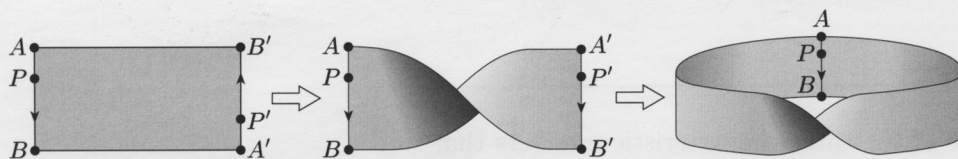


Figure 1.26 Making a Möbius band from a rectangle

Torus

Let us now return to the cylinder we obtained in Figure 1.25. What happens if we bend it around and glue the two ends? Now, continuing with our idea of identifying edges of our original rectangle, we want to bend the cylinder round and glue its ends in such a way that the points A and B are identified. Furthermore, bending as shown in Figure 1.27 corresponds to identifying the edges AA' and BB' of the original rectangle *in the same direction*, as indicated in the figure. We obtain a *torus*.

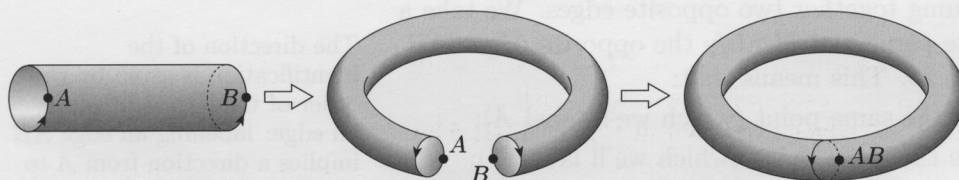


Figure 1.27 Making a torus from a cylinder

In terms of our original rectangle, the whole construction process corresponds to identifying opposite edges in pairs in the directions indicated by the arrowheads in Figure 1.28. For ease of reference, we label each pair of edges that is to be identified with the same letter, a for the first pair to be identified and b for the second pair.

Here, we are identifying the two ends of the rectangle rather than the top and bottom.

Again, the labelling AB implies a direction from A to B , and similarly $A'B'$ implies a direction from A' to B' .

We often omit the vertex labels from polygons used to construct surfaces by identification and simply use these edge labels.

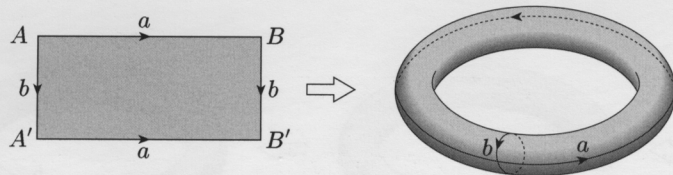


Figure 1.28 Making a torus from a rectangle

Klein bottle

There are two other surfaces that can be obtained by identifying both pairs of opposite edges of a rectangle. In one of these, shown in Figure 1.29, we first identify the edges AB and $A'B'$, labelled a , in the direction shown by the arrowheads. This gives us a cylinder, as before. We then try to identify the cylinder ends, labelled b , in *opposite* directions, as indicated by the arrowheads — this is equivalent to identifying the edges $A'A$ and BB' , labelled b , of the original rectangle. Unfortunately, this cannot be done in three dimensions, and so is hard to visualize: if we try to link the ends, the cylinder has to ‘pass through itself’. So, although we have a surface, it is not a surface in space. It is called a **Klein bottle**.

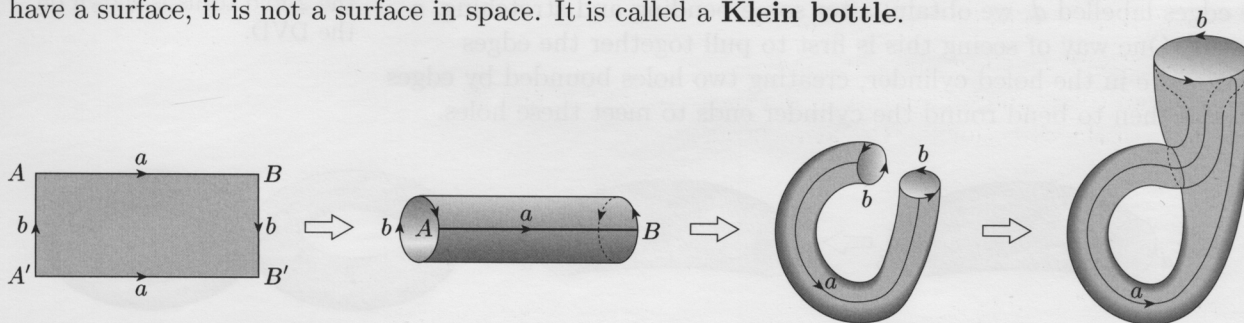


Figure 1.29 Making a Klein bottle from a rectangle

Projective plane

The final surface that can be obtained by identifying edges of a rectangle is even more complicated. Again it cannot be constructed in three dimensions, so is not a surface in space and is hard to visualize. This time we identify *both* pairs of opposite edges of the rectangle in opposite directions, as indicated by the arrowheads in Figure 1.30: we identify the edges AB and $B'A'$ (labelled a) and the edges $A'A$ and BB' (labelled b). The resulting surface is called the **projective plane**.

The existence of the Klein bottle and the projective plane, which are not surfaces in space, is one reason why mathematicians cannot restrict their study of surfaces to surfaces in space.

Torus with 1 hole

We need not restrict ourselves to rectangles: we can also build surfaces by identifying edges of other polygons. For example, if we start with a pentagon and identify two pairs of its edges as shown in Figure 1.31, what do we get? Identifying the edges labelled a and c in the directions indicated, we obtain a cylinder with a hole; the boundary of the hole is the edge labelled c . If we now identify the edges labelled b (now at each end of the cylinder) in the directions indicated, after some bending and stretching we obtain a torus with a hole.

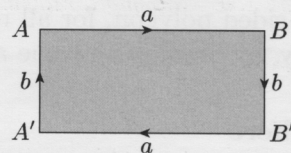


Figure 1.30 Projective plane rectangle

We discuss the projective plane in Section 2.

From now on, we shall usually put a direction on *every* edge (such as edge c), even if it is not to be identified. For such an edge, we can choose an arbitrary direction.

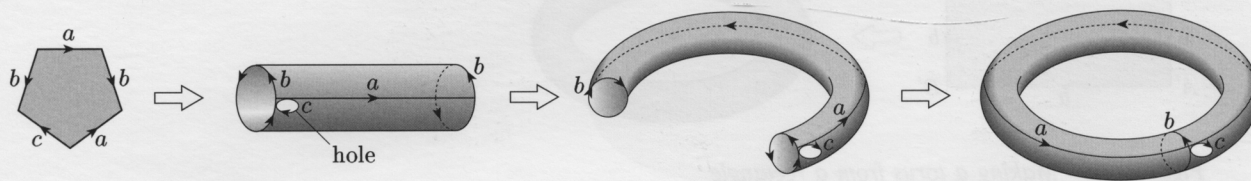


Figure 1.31 Making a torus with 1 hole from a pentagon

2-fold torus

As the polygons become more complicated, so the identifications become difficult to visualize. For example, what happens if we try to identify the edges of an *octagon* in pairs, as indicated by the edge labels and arrowheads in Figure 1.32? Figure 1.33 shows that identifying, in the directions indicated, the edges labelled *a* and the edges labelled *c* results, after some bending and stretching, in a cylinder bounded by edges labelled *b* and *d*. If we now identify, in the directions indicated, the edges labelled *b* and the edges labelled *d*, we obtain, after some bending and stretching, a 2-fold torus. One way of seeing this is first to pull together the edges labelled *a* and *c* in the holed cylinder, creating two holes bounded by edges *b* and *d*, and then to bend round the cylinder ends to meet these holes.

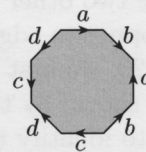


Figure 1.32

Another way of constructing the 2-fold torus is shown on the DVD.

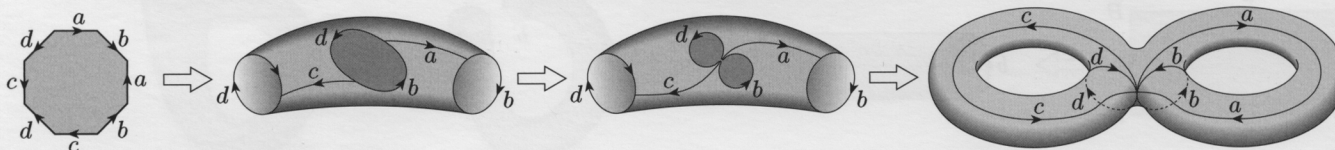


Figure 1.33 Making a 2-fold torus from an octagon

In a similar way, one could obtain a 3-fold torus by identifying the twelve edges of a dodecagon in pairs, and in general an n -fold torus from a $4n$ -sided polygon, for all natural numbers n . Instead we'll present a better way, by developing some algebraic techniques (in *Unit B2*).

Sphere

Surfaces can be constructed in a similar way from plane figures other than polygons. For example, starting with a disc, we can fold the left-hand half over onto the right-hand half, and identify the edges labelled *a*, as shown in Figure 1.34; this is rather like zipping up a purse, or 'crimping' a Cornish pasty. We can then stretch and inflate the object so obtained until we obtain a sphere.

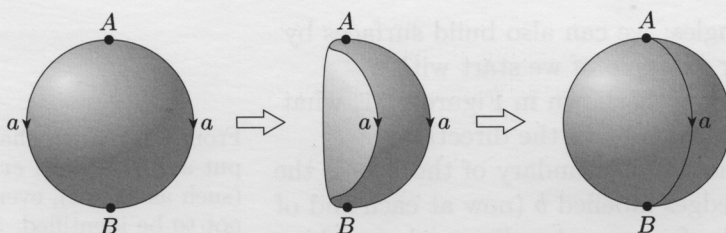


Figure 1.34 Making a sphere from a disc

We can even start with more than one plane figure. For example, a sphere can be formed from two discs by identifying edges as shown in Figure 1.35 and then inflating.

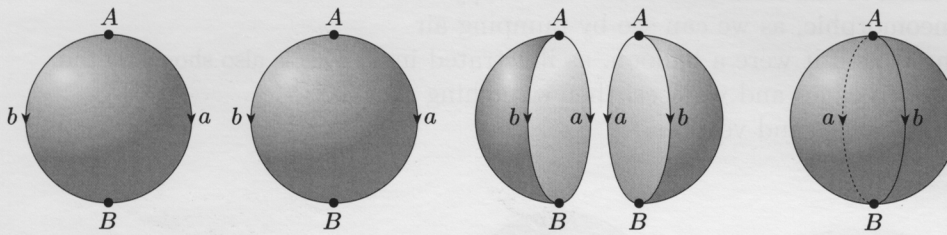


Figure 1.35 Making a sphere from two discs

Problem 1.9

Which surfaces are obtained by identifying edges (in the given directions) in each of the triangles in Figure 1.36?

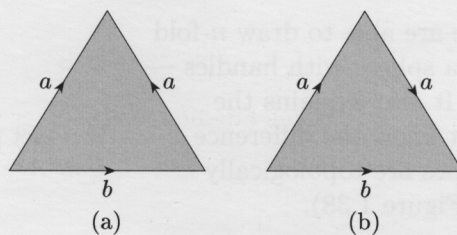


Figure 1.36

Hint In (b), you might consider cutting the triangle before identifying the edges, and then repairing the cut.

Throughout this subsection, our emphasis has been on describing the edge identifications geometrically. A more rigorous treatment, involving the so-called ‘identification topology’, is given in Section 4.

1.4 Homeomorphic surfaces

As we stated in the Introduction, our aim is to classify surfaces up to homeomorphism. So it is worthwhile spending a little time examining what sorts of transformations of surfaces are homeomorphisms. We shall restrict the description to surfaces in space, as these are easier to deal with, though the result at the end of this subsection applies to *all* surfaces.

Recall that a homeomorphism between two topological spaces (such as surfaces in space) is a bijection with the property that both f and its inverse f^{-1} preserve open sets.

Roughly speaking, homeomorphisms between surfaces in space are of two kinds. One kind includes any transformation of a surface in space that can be achieved by bending, stretching, squeezing or shrinking the surface. These maps are certainly one-one and onto, and (intuitively, at least) seem to be continuous — and thus are bijections. They would also seem to preserve open sets, if we remember that the topology of a surface in space is the subspace topology inherited from \mathbb{R}^3 . So, essentially, we can achieve a homeomorphism by treating a surface in space as if it were made from a sheet of rubber and then bending, stretching, squeezing or shrinking it.

Topology is sometimes called *rubber-sheet geometry*.

For example, the sphere in its usual round form and in various lumpy forms (such as a cube) are homeomorphic, as we can see by pumping air into the lumpy one and inflating it as if it were a balloon, as illustrated in Figure 1.37. If the surface has sharp edges and vertices, then stretching and smoothing it can remove these edges and vertices.

This is also shown on the DVD.

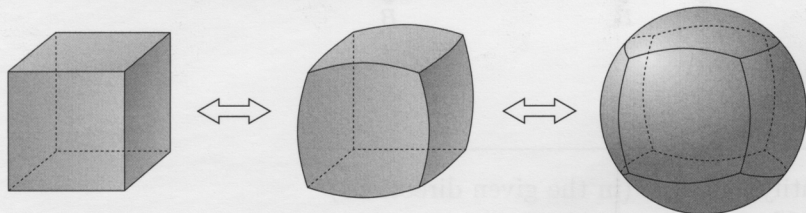


Figure 1.37 Inflating a cube

This sort of homeomorphism also explains why we are able to draw n -fold toruses either in doughnut form or in the form of a sphere with handles — both representations are topologically equivalent. It also explains the description of a topologist as someone who doesn't know the difference between a doughnut and a teacup, since both shapes are topologically equivalent under a homeomorphism of this kind (Figure 1.38).

You met this description in Unit A3.

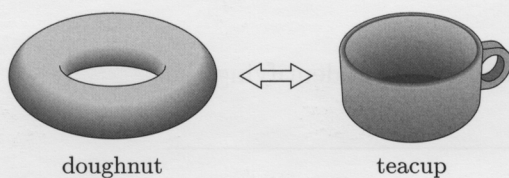


Figure 1.38

As an example of the other kind of homeomorphism, consider the torus in its usual form, sometimes referred to as its standard form, and in a knotted form, as shown in Figure 1.39.

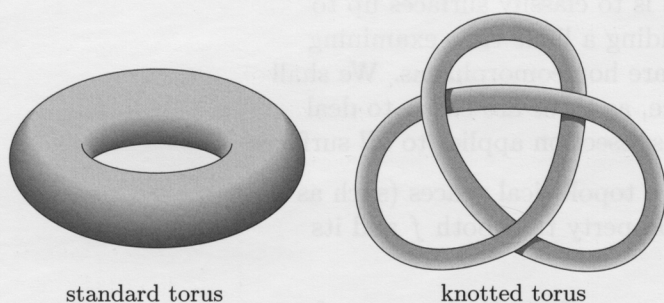


Figure 1.39

It is impossible to 'unknot' the knotted torus just by bending, stretching, squeezing or shrinking it — but we can cut it and open it out, rearrange it as a cut torus, and then glue it back together, as shown in Figure 1.40.

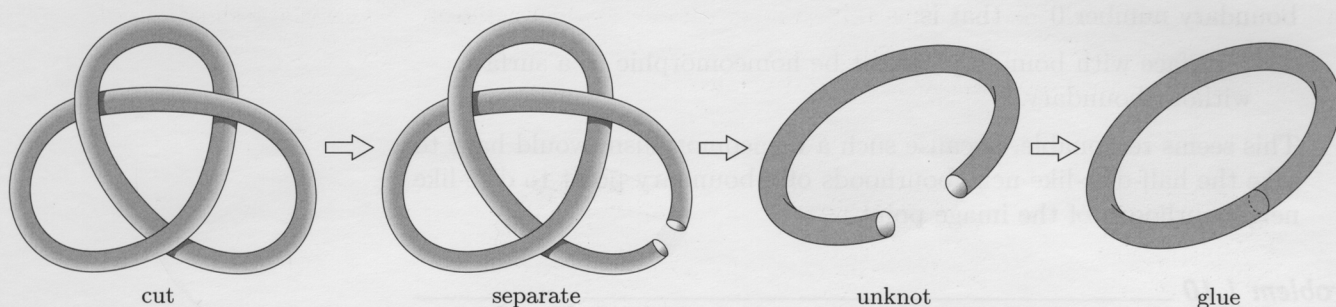


Figure 1.40

Provided that we are able to join the points on either side of the cut exactly as they were beforehand, the whole operation defines a homeomorphism between the two surfaces. If this seems surprising, look at the neighbourhood of a point on the cut, as enlarged in Figure 1.41: clearly, open discs are mapped to open discs without any change.

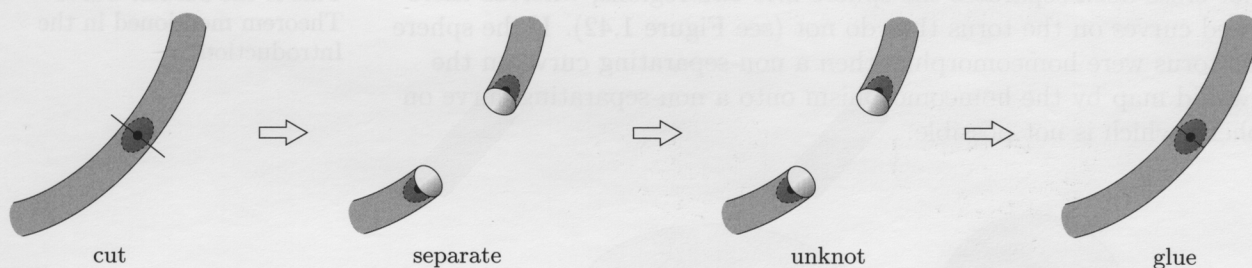


Figure 1.41

To sum up, two surfaces in space are homeomorphic if we can bend, stretch, squeeze or shrink one into the other and/or if we can cut one and then, after some bending, stretching, squeezing or shrinking, glue it back together (making sure to join the points on either side of the cut exactly as beforehand) to form the other. We thus have a practical and intuitive way of demonstrating that surfaces in space are homeomorphic.

There are also practical ways of showing that two surfaces in space are not homeomorphic. One is to consider the boundary number of the surface.

Theorem 1.2

Two surfaces that are homeomorphic have the same boundary number.

We omit the proof, which is beyond the scope of this unit.

Remarks

- (i) This theorem applies to *all* surfaces and not just to surfaces in space.
- (ii) This theorem tells us that the boundary number is a **topological invariant** for surfaces, i.e. a property that is invariant under homeomorphisms.
- (iii) It follows from the theorem that two surfaces with different boundary numbers cannot be homeomorphic. It does not follow that two surfaces with the same boundary number *are* homeomorphic — for example, a sphere and a torus both have boundary number 0, but are not homeomorphic (as we demonstrate below).

- (iv) Another consequence of the theorem is that a surface with non-zero boundary number cannot be homeomorphic to a surface with boundary number 0 — that is:

a surface with boundary cannot be homeomorphic to a surface without boundary.

This seems reasonable, because such a homeomorphism would have to take the half-disc-like neighbourhoods of a boundary point to disc-like neighbourhoods of the image point.

Problem 1.10

Use boundary numbers to show that a Möbius band and a 2-fold torus are not homeomorphic.

For surfaces in space with the same boundary number, there is another practical test that can sometimes be used to confirm that they are not homeomorphic. It requires the consideration of closed curves on the surface. For example, consider a sphere and a torus, both of which have boundary number 0. Observe that any closed curve on the sphere that does not cross itself separates the sphere into two regions, whereas there are closed curves on the torus that do not (see Figure 1.42). If the sphere and the torus were homeomorphic, then a non-separating curve on the torus would map by the homeomorphism onto a non-separating curve on the sphere, which is not possible.

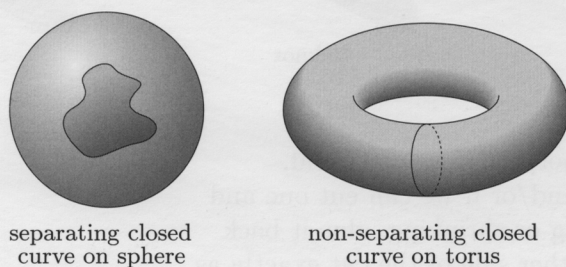


Figure 1.42

This is the Jordan Curve Theorem mentioned in the Introduction.

1.5 Defining surfaces

In Subsection 1.1 we gave a provisional definition of a surface. The aim of this subsection is to formalize that definition. To do that, we need to specify three further requirements of a candidate topological space, beyond those given in the provisional definition.

The first requirement is that the surface should be in 'just one piece'. We can ensure this by requiring the surface to be **path-connected**: this means that any two points P and Q on the surface can be joined by a curve that lies entirely in the surface. In mathematical terms, this means that there is a continuous map f from the interval $[0, 1]$ to the surface, such that $f(0) = P$ and $f(1) = Q$ (see Figure 1.43). (Path-connectedness is discussed in detail in *Unit C1*.)

The second requirement is a technical one that is needed to eliminate certain awkward cases. We require a surface to be a **Hausdorff space**: this means that, given any pair of distinct points a and b in it, there are *disjoint* open sets U and V , one containing a and the other containing b , as shown in Figure 1.44. (We can think of the points a and b as being

We mention these requirements briefly here, but will not study them in detail until Block C.

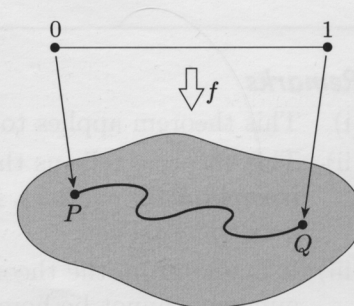


Figure 1.43

separated, or ‘housed orff’, from each other by open sets!) Hausdorff spaces are named after their originator Felix Hausdorff (1868–1942) and are discussed in detail in *Unit C2*.

Not all topological spaces are Hausdorff: for example, the topological space consisting of the set $X = \{a, b\}$ with the indiscrete topology is not Hausdorff, since the only non-empty open set is X itself, which does not separate a from b . However, \mathbb{R}^3 with the Euclidean topology is Hausdorff, as is any subset of it with the induced topology: so surfaces in space are Hausdorff. Thus, for a topological space to be a surface in space, it must be Hausdorff.

We shall see in Block C that being path-connected and being Hausdorff are both topological invariants: thus, if two topological spaces are homeomorphic and one of them is path-connected, then so is the other; and, if one of them is Hausdorff, then so is the other.

Our final requirement for a surface is that it be *compact*. Such surfaces can be defined in a number of ways. For our purposes in this block, we shall think of compact surfaces as follows.

Definition

A **compact surface** is a surface that can be obtained from a polygon (or a finite number of polygons) by identifying edges.

For example, the surfaces we constructed in Subsection 1.3 — cylinder, Möbius band, torus, Klein bottle, projective plane, torus with 1 hole, 2-fold torus and sphere — are all compact surfaces. Examples of surfaces that are not compact are a cylinder without its bounding circles and the entire plane. Compactness is also a topological invariant, as we shall see in *Unit C2*, so any topological space that is homeomorphic to a compact surface is also a compact surface.

The importance of compact surfaces is that they are precisely the surfaces that we can classify by means of the Classification Theorem. This is because they are the surfaces for which we can guarantee that the numbers used to classify surfaces are finite. For other types of surface, it can be very difficult to decide when two surfaces are homeomorphic.

Including the three new requirements that a surface be path-connected, Hausdorff and compact, we are able to turn our provisional definition of a surface into a formal definition.

Definition

A **surface** is a compact path-connected Hausdorff topological space (X, \mathcal{T}) with the property that, given any point $x \in X$, there is an open set U containing x such that U is homeomorphic either to an open disc in \mathbb{R}^2 with the Euclidean topology or to an open half-disc in the upper half-plane with the subspace topology inherited from the Euclidean topology on \mathbb{R}^2 .

A **surface in space** is a surface (X, \mathcal{T}_X) where X is a subset of \mathbb{R}^3 and \mathcal{T}_X is the subspace topology on X inherited from the Euclidean topology \mathcal{T} on \mathbb{R}^3 .

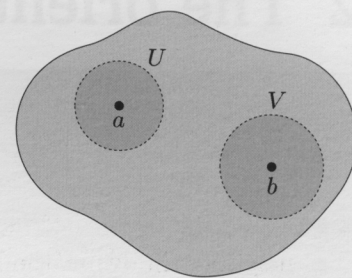


Figure 1.44

We shall discuss general compact topological spaces from several points of view in *Unit C2*.

We state the Classification Theorem in Section 3, and prove it in *Unit B3*.

It is this definition of a surface that we shall use throughout Block B.

2 The orientability of surfaces

After working through this section, you should be able to:

- ▶ explain the terms *orientable* and *non-orientable* with regard to surfaces;
- ▶ describe the *projective plane*, and explain why it is non-orientable;
- ▶ outline the connections between the projective plane and the Möbius band.

In this section we study the orientability of surfaces from an informal point of view. In particular, we take a detailed look at the projective plane and its properties. We start by examining some surfaces that resemble a Möbius band.

2.1 Surfaces with twists

A cylinder or a Möbius band can be formed by gluing together the ends of a rectangular strip or band of paper either with or without twisting the paper before gluing. Does adding further twists to the band before gluing provide any more examples of surfaces?

We say that the band shown on the left of Figure 2.1 has a *half-twist* — and going from left to right it is a *clockwise* half-twist — while the band on the right has two clockwise half-twists, making a *whole twist*.

A half-twist is through π radians; a whole twist is through 2π radians.

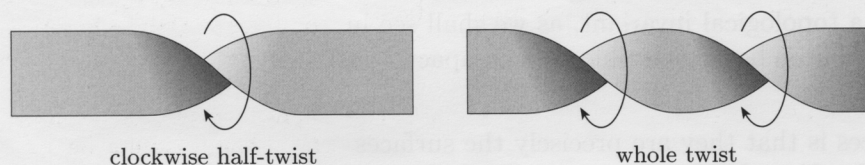


Figure 2.1

Clearly, the twists must be in the same direction (clockwise or anticlockwise), or they will cancel each other out, as Figure 2.2 shows.

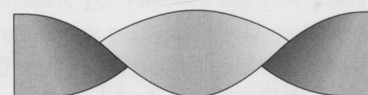


Figure 2.2 Two half-twists that cancel out

If we join up the ends of a band with a half-twist, we get a Möbius band, as shown in Figure 2.3(a). But what do we get if we join up the ends of a band with a whole twist, as in Figure 2.3(b)?

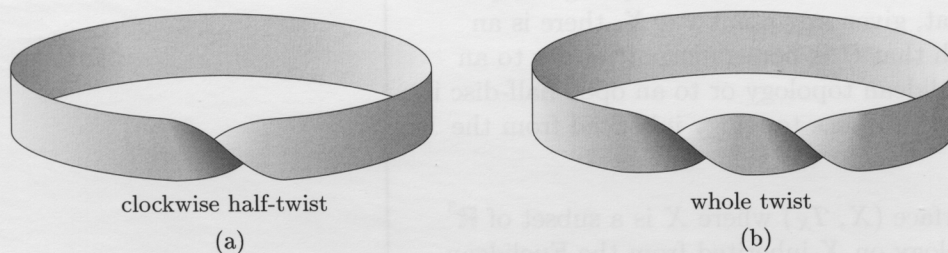


Figure 2.3

To answer this question, we cut the band, as shown in Figure 2.4. We can then undo the whole twist, glue the cut edges back together (making sure

that the two edges are in the same direction, as before), and obtain a cylinder.

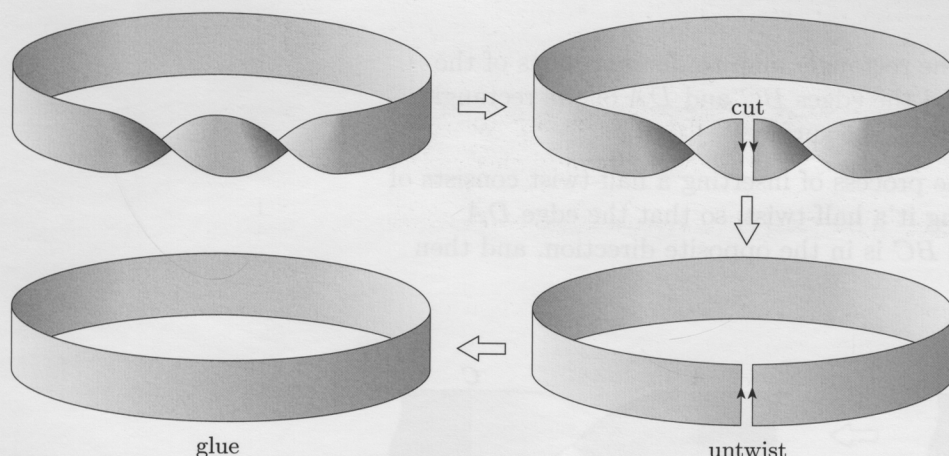


Figure 2.4

This cutting and gluing process is an example of a homeomorphism, as we saw in Subsection 1.4. Thus:

a band with a whole twist is homeomorphic to a band with no twists.

Clearly, using the method outlined above, we can undo any number of whole twists, so that a band with any number of whole twists is homeomorphic to a band with no twists: a cylinder. However, we cannot undo a half-twist. Thus, for example, if we have a band with three whole twists and an extra half-twist, then we can untwist the three whole twists but we cannot do anything with the half-twist, so we have a Möbius band.

We can conclude that, topologically, there are just two types of band:

- bands with an even number of half-twists, which are homeomorphic to a cylinder;
- bands with an odd number of half-twists, which are homeomorphic to a Möbius band.

This categorization of bands will prove useful when we discuss orientability in Subsection 2.2.

Problem 2.1

Write down the boundary number β of a band with:

- (a) three half-twists; (b) four half-twists.

The intuitive difference we feel between a band with a number of twists and a band with no twists (cylinder) or one twist (Möbius band) is captured by a theorem which states that the homeomorphisms between such bands cannot be extended to homeomorphisms of \mathbb{R}^3 .

Inserting half-twists

We can insert half-twists into a paper surface whenever a piece of the surface is homeomorphic to a rectangle $ABCD$ with the following properties:

the edges AB and CD of the rectangle map to distinct parts of the boundary of the surface, and the edges BC and DA of the rectangle map to non-boundary points of the surface.

As illustrated in Figure 2.5, the process of inserting a half-twist consists of cutting out the rectangle, giving it a half-twist, so that the edge DA appears as it was but the edge BC is in the opposite direction, and then gluing it back.

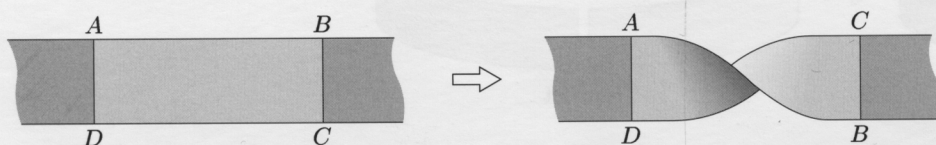


Figure 2.5 Inserting a twist

Since inserting a half-twist is not a homeomorphism, this process allows us to create new surfaces.

As well as creating a new surface, inserting a half-twist can change the boundary number. For example, Figure 2.6(a) shows a paper surface with boundary number $\beta = 3$. If we give the marked rectangle a half-twist, as shown on the right, then we obtain a surface with boundary number $\beta = 2$.

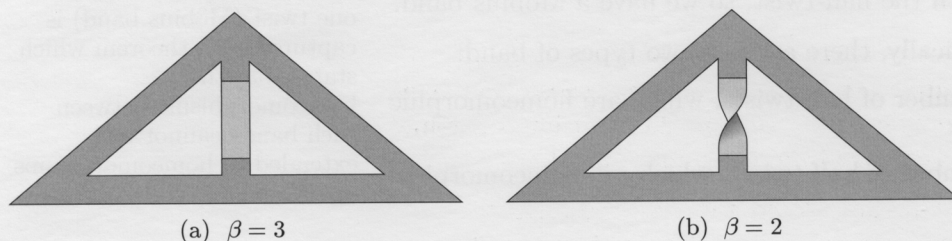


Figure 2.6

Problem 2.2

Starting with the paper surface shown in Figure 2.6(b), insert a half-twist to produce a surface with boundary number $\beta = 1$.

2.2 Orientability

The idea of orientability is another fundamental concept that we need for the study of surfaces. To illustrate the underlying idea, we consider two familiar surfaces — a cylinder and a Möbius band.

We can distinguish between a cylinder and a Möbius band by noticing that, every cylinder has an ‘inside’ and an ‘outside’, as shown in Figure 2.7 — and we can paint one red and the other blue — but, if we try to paint a Möbius band in two colours, we fail — it has just one ‘side’. Equivalently, an ant crawling on a Möbius band can visit all points of the surface without crossing the boundary edge, whereas an ant on a cylinder can visit only the inside or the outside without crossing the boundary edges.

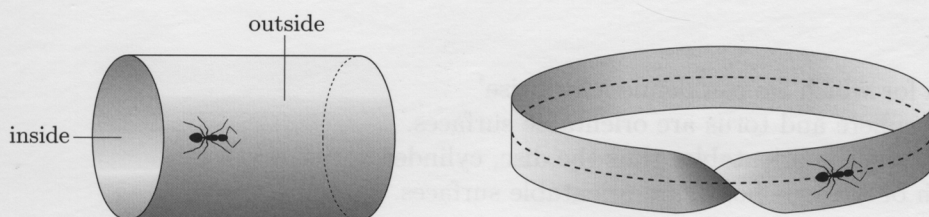


Figure 2.7

Two-sided surfaces in space, such as a cylinder, are examples of *orientable* surfaces, whereas one-sided surfaces in space, such as a Möbius band, are examples of *non-orientable* surfaces.

An **orientable surface** is one where consistent ‘orientation’ can be assigned over the entire surface; a **non-orientable surface** is one where this cannot be done.

To see what we mean by this, let us take a cylinder, and place on its surface a clock-face whose minute hand is fixed at 12 and whose hour hand is fixed at 3, as shown in Figure 2.8(a): notice that the hour hand is found by going through $\pi/2$ radians *clockwise* from the minute hand.

If we now slide the clock-face along any closed curve on the surface, keeping the hour hand pointing along the curve, then the hour hand is always $\pi/2$ radians clockwise from the minute hand: for example, if the hour hand moves along the curve C shown in Figure 2.8(b), then the clock-face continues to show 3 o’clock. This means that we can consistently define ‘clockwise’ at any point of the surface: the hour hand is always $\pi/2$ radians *clockwise* from the minute hand.

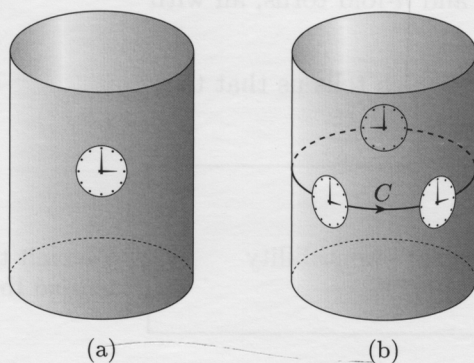


Figure 2.8

A similar situation occurs if we slide a clock-face along any closed curve C on the surface of a sphere or a torus, as shown in Figure 2.9: in each case the hour hand is always $\pi/2$ radians clockwise from the minute hand, and so we can again define 'clockwise' consistently.

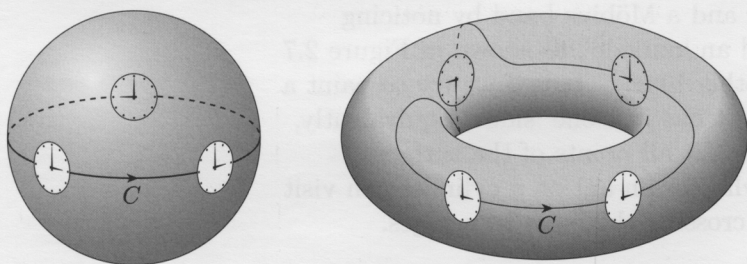


Figure 2.9

Orientable surfaces are surfaces for which we can define 'clockwise' consistently: thus, the cylinder, sphere and torus are orientable surfaces. In fact, any two-sided surface in space is orientable: thus the disc, cylinder, sphere and n -fold torus, all with or without holes, are orientable surfaces.

Let us now consider a clock-face on a Möbius band, and let us slide it around the central curve C , as shown in Figure 2.10. Starting from any point on the curve, the twist in the surface causes the clock-face to be upside-down when it first returns to that point: the hour hand still points to 3, but the minute hand now points downwards. Thus, the hour hand is no longer $\pi/2$ radians clockwise from the minute hand. This means that we cannot consistently define 'clockwise' at every point of a Möbius band. Thus, the Möbius band is a non-orientable surface. In fact, any one-sided surface in space is non-orientable.

Just as we use the boundary number β to distinguish between surfaces without boundary ($\beta = 0$) and surfaces with boundary ($\beta > 0$), so we introduce a number that distinguishes between orientable and non-orientable surfaces. This is another of the three characteristic numbers we use to classify compact surfaces.

Definition

The **orientability number** ω of a surface is:

- 0 if the surface is orientable;
- 1 if the surface is non-orientable.

For example, $\omega = 0$ for the disc, cylinder, sphere, and n -fold torus, all with or without holes, and $\omega = 1$ for the Möbius band.

As you might expect, we have the following result, which tells us that the orientability number is a topological invariant.

Theorem 2.1

Two surfaces that are homeomorphic have the same orientability number.

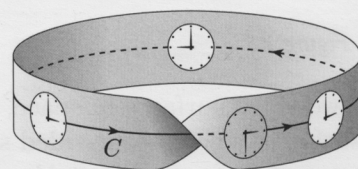


Figure 2.10

As we shall soon see, the Klein bottle and projective plane are also non-orientable.

ω is the Greek letter omega.

We omit the proof, which is beyond the scope of this unit.

Problem 2.3

Write down the orientability number ω of each of the following surfaces:

- (a) the 3-fold torus; (b) a band with three half-twists.

One way of determining the orientability number of a given surface is to ascertain whether the surface contains a Möbius band.

Theorem 2.2

A surface that contains no Möbius band is orientable ($\omega = 0$).
A surface that contains a Möbius band is non-orientable ($\omega = 1$).

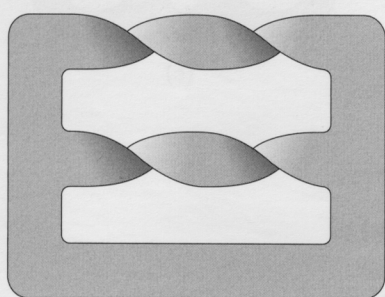
We omit the proof, which is beyond the scope of this unit.

Remarks

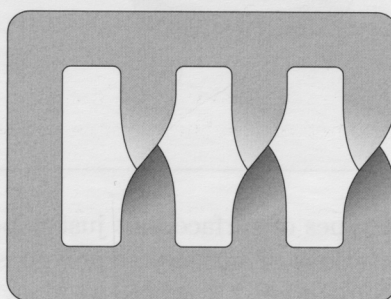
- (i) By 'contains', we mean that we can find part of the surface that is homeomorphic to a Möbius band. The edge of the Möbius band does not need to correspond to an edge at the surface, so that a surface without boundary can be non-orientable (as we shall shortly see).
- (ii) When seeking Möbius bands in a surface, it can be helpful to look at all possible closed curves on the surface, and thicken these into bands.
- (iii) Remember, from Subsection 2.1, that all bands with an odd number of half-twists (in the same direction) are homeomorphic to Möbius bands.

Worked problem 2.1

Determine whether each of the paper surfaces in Figure 2.11 is orientable.



(a)



(b)

Figure 2.11

Solution

- (a) The top and middle of this surface have two half-twists each. These can be cut and glued back together to give the surface in Figure 2.12.

This surface is a disc with 2 holes, which is orientable ($\omega = 0$).

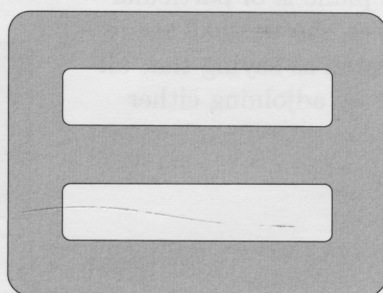


Figure 2.12

- (b) This surface contains several Möbius bands, two of which correspond to the closed curves shown in Figure 2.13.

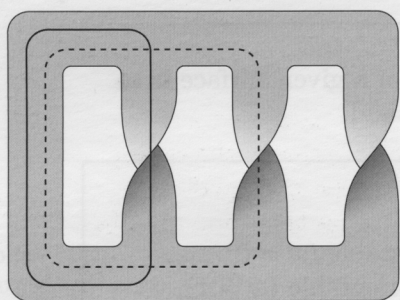


Figure 2.13

The surface is therefore non-orientable ($\omega = 1$). ■

Problem 2.4

Determine whether each of the paper surfaces in Figure 2.14 is orientable.

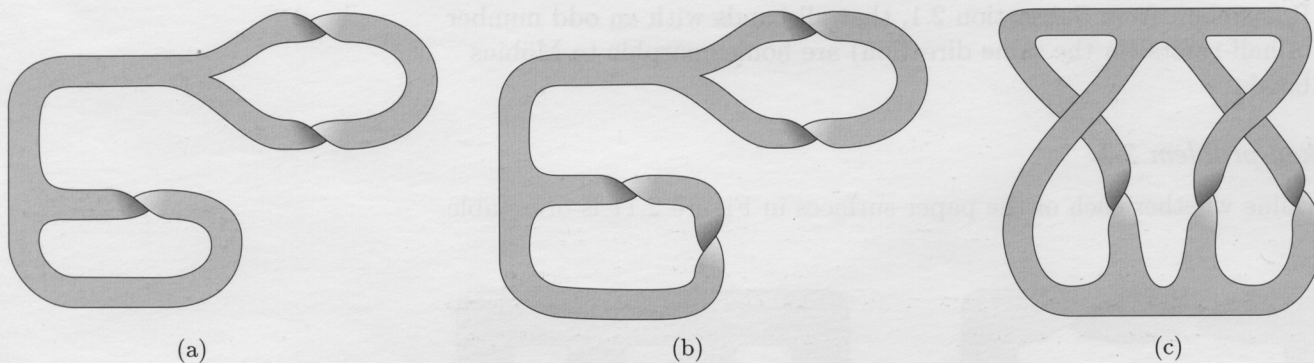


Figure 2.14

Theorem 2.2 can be applied to all types of surface, not just paper surfaces. It can even be applied to representations of surfaces as polygons with edge identifications, as we shall see in Subsection 2.3, where we show that the projective plane and the Klein bottle are non-orientable.

2.3 The projective plane

We now consider one of the most important non-orientable surfaces — the *projective plane*. In Section 1 we introduced it as the surface obtained from a rectangle by identifying each pair of opposite edges in opposite directions, as shown in Figure 2.15. The projective plane is of particular importance in relation to the classification of surfaces. As we shall see in *Unit B3*, the Classification Theorem can be interpreted as saying that all surfaces without boundary can be made from a sphere by adjoining either toruses or projective planes.

It is sometimes called the *real projective plane*.

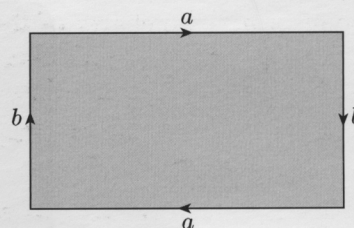


Figure 2.15

You may already be familiar with a geometric way of viewing the projective plane. This gives a completely different way of thinking about it: indeed, it is not even immediately obvious from the geometric approach why it is a surface!

The geometric approach is to define the projective plane as the set of all infinite lines through the origin in Euclidean three-dimensional space: this means that each infinite line through the origin is to be regarded as a 'point' of the projective plane. To do this, we take \mathbb{R}^3 , remove the origin $(0, 0, 0)$, and consider any two of the remaining points (x, y, z) and (x', y', z') of \mathbb{R}^3 to be *equivalent* if and only if they lie on the same infinite line through the origin — that is, $x' = kx$, $y' = ky$, $z' = kz$, for some non-zero real number k . For example, the points $(1, 2, 3)$, $(2, 4, 6)$ and $(-10, -20, -30)$ of \mathbb{R}^3 are all equivalent points, as Figure 2.16 illustrates. Each set of equivalent points, together with the point $(0, 0, 0)$, comprises an infinite line through the origin and is thus a 'point' of the projective plane, which we refer to as a **projective point**. We denote a projective point using square brackets around the coordinates of a representative point on the infinite line: thus the projective point containing the points $(1, 2, 3)$, $(2, 4, 6)$ and $(-10, -20, -30)$ may be written as $[1, 2, 3]$ or $[2, 4, 6]$, or indeed as $[k, 2k, 3k]$ for any non-zero number k .

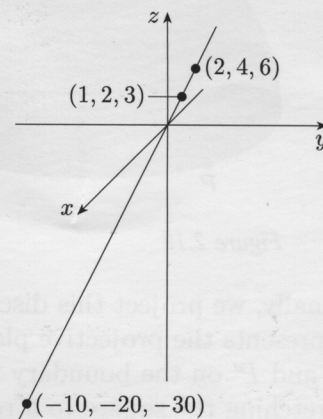


Figure 2.16

Problem 2.5

Which of the following points (x, y, z) of \mathbb{R}^3 correspond to the same projective points $[x, y, z]$?

$(3, -1, 2)$, $(4, -2, 6)$, $(-9, 3, -6)$, $(12, -6, -9)$, $(15, -5, 10)$, $(-18, 9, -27)$.

This definition of the projective plane is no more helpful than the definition in terms of a rectangle with edge identifications as a means of picturing what the projective plane really looks like. It does, however, lead us towards a description of the projective plane in which each projective point looks like a point of \mathbb{R}^3 . We begin by enclosing the origin within the sphere of unit radius, as shown in Figure 2.17. Each infinite line through the origin gives rise to a diameter of this sphere which meets the sphere in exactly two antipodal points, (p, q, r) and $(-p, -q, -r)$. So we can think of the projective plane as the set of all these *pairs of antipodal points* on the unit sphere, and the projective point $[p, q, r]$ now corresponds to the *two* points (p, q, r) and $(-p, -q, -r)$ of \mathbb{R}^3 on the sphere. This gives a natural 'two-to-one' map from the sphere to the projective plane which sends the two points (p, q, r) and $(-p, -q, -r)$ to the projective point $[p, q, r]$.

So we can regard the projective plane as made up of pairs of points on the sphere. If we could systematically select a *single* point from each of these pairs, then we would have an even better picture of the projective plane. Unfortunately, there is no natural way of thinking of *exactly half* of the points on a sphere. We can certainly think of the points $[p, q, r]$ with $r \neq 0$ as lying on the *upper hemisphere*: we simply associate with each projective point $[p, q, r]$ whichever of the two points (p, q, r) and $(-p, -q, -r)$ on the sphere has its third coordinate positive. But if $r = 0$, then we have a problem: the projective point $[p, q, r]$ is associated with both the points $(p, q, 0)$ and $(-p, -q, -0)$ of the sphere, and there is no satisfactory way of choosing just one of them.

Instead, we proceed as follows. We first represent the projective plane as the upper hemisphere with antipodal pairs of points on the equator ($r = 0$) identified. We then map this space on to the unit disc by projecting vertically downwards, using the map $(x, y, z) \mapsto (x, y, 0)$, as shown in

This sphere has equation $x^2 + y^2 + z^2 = 1$.

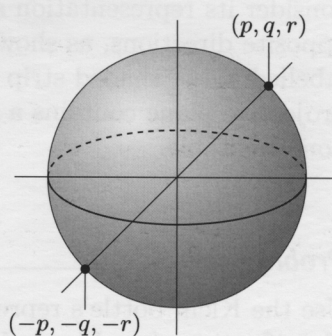


Figure 2.17

Figure 2.18. We thereby represent the projective plane as a disc with antipodal points P and P' on the boundary identified.

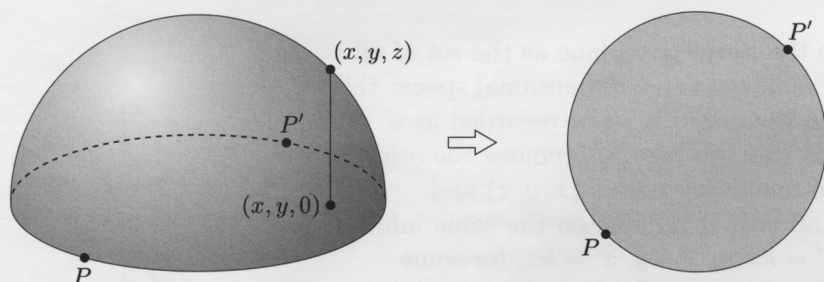


Figure 2.18

Finally, we project this disc radially onto a circumscribing square. This represents the projective plane as a square with diagonally opposite points P and P' on the boundary identified, as shown in Figure 2.19. By stretching the square to a rectangle, we obtain our earlier description of the projective plane as a rectangle with opposite edges identified in opposite directions.

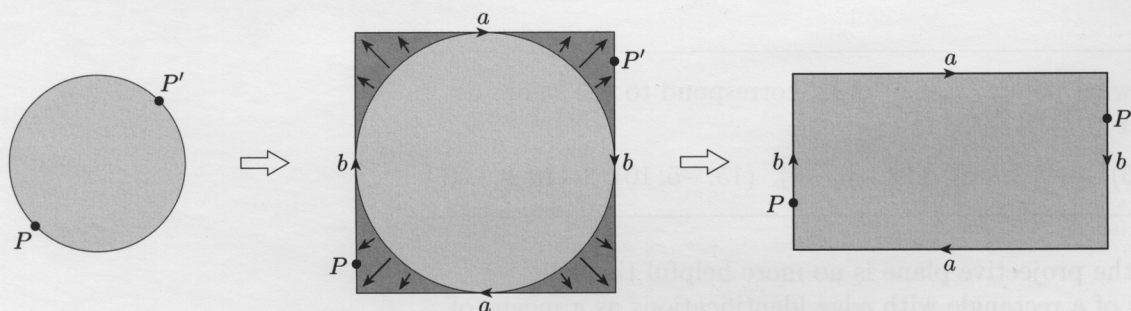


Figure 2.19

In Subsection 2.2 we asserted that a surface is non-orientable if it contains a Möbius band. To show that the projective plane is non-orientable, we consider its representation as a rectangle with opposite edges identified in opposite directions, as shown in Figure 2.20. When we identify the edges labelled b , the shaded strip becomes a Möbius band. This shows that the projective plane contains a Möbius band, and therefore by Theorem 2.2 is non-orientable.

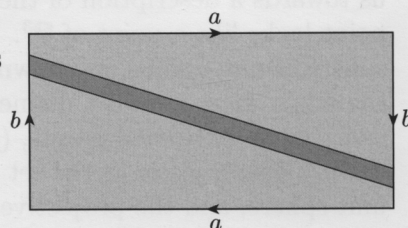


Figure 2.20 A Möbius band in the projective plane

Problem 2.6

Use the Klein bottle's representation as a rectangle with edge identifications (see Figure 1.29) to show that it is non-orientable.

We can obtain the same result using the representation of the projective plane as an upper hemisphere with antipodal points identified. To do so, we draw half of a great circle Γ on it through two antipodal points A, A' and the north pole. We then 'fatten' this line to a narrow strip that meets the equator at the pairs of antipodal points B, B' and C, C' , as shown in Figure 2.21. The boundary of this strip around Γ is a single curve, and the shaded strip is a Möbius band, showing that the surface is non-orientable. Either way, the presence of a Möbius band shows that the projective plane is non-orientable. This result is so important that we state it as a theorem.

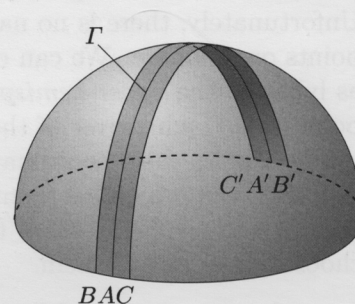


Figure 2.21

Theorem 2.3

The projective plane is non-orientable.

If we remove a Möbius band from the projective plane, what remains? Using the rectangle representation, we first obtain two pieces, as shown in Figure 2.22(a), with the edges labelled a , b and c to be identified in pairs; the edges x and y form the boundary of the removed Möbius band. Turning each piece over and identifying the edges labelled a , we obtain Figure 2.22(b). Identifying the edges labelled b and then the edges labelled c , we obtain a purse-like object whose boundary is the boundary of the Möbius band, as shown in Figure 2.22(c). Finally, if we open the purse and flatten it out, we obtain a disc whose boundary is the boundary of the Möbius band, as shown in Figure 2.22(d).

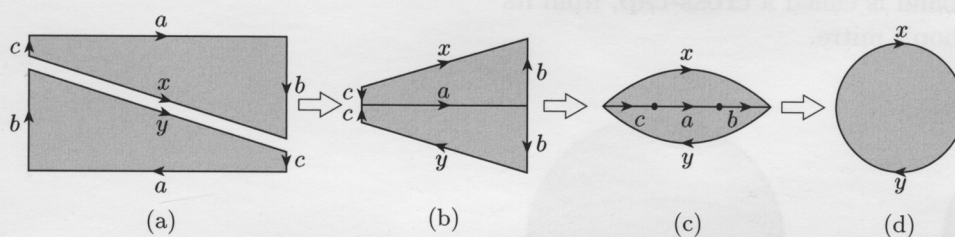


Figure 2.22

Thus, removing a Möbius band from the projective plane yields a disc whose boundary is the boundary of the Möbius band. Gluing the Möbius band back again leads to the following surprising conclusion.

Theorem 2.4

The projective plane is obtained by gluing a disc and a Möbius band along their boundaries.

We can use this result to provide us with an alternative picture of a Möbius band. We begin with a representation of the projective plane as a hemisphere. However, instead of taking the upper hemisphere as before, it is now more convenient to represent it as the lower hemisphere, but again with antipodal points on the equator identified, as shown in Figure 2.23(a). We then distort the hemisphere slightly, giving the equator a 'flower' shape, as shown in Figure 2.23(b): this makes it easier for us to see what is to be identified with what. We see that the edge AB is to be identified with the edge $A'B'$, and that the edge BA' is to be identified with the edge $B'A$. We can perform either one of these two identifications without too much difficulty after some judicious stretching of the surface; but, when we try to perform the other, we need to let the surface 'intersect itself' along a line, the vertical line in Figure 2.23(c).

Thus Figure 2.23(c) gives us a representation of the projective plane in which no identifications need to be made, but it is not a representation in \mathbb{R}^3 since the surface pictured intersects itself.

We saw a similar picture of the Klein bottle, as a surface that intersects itself, in Figure 1.29.

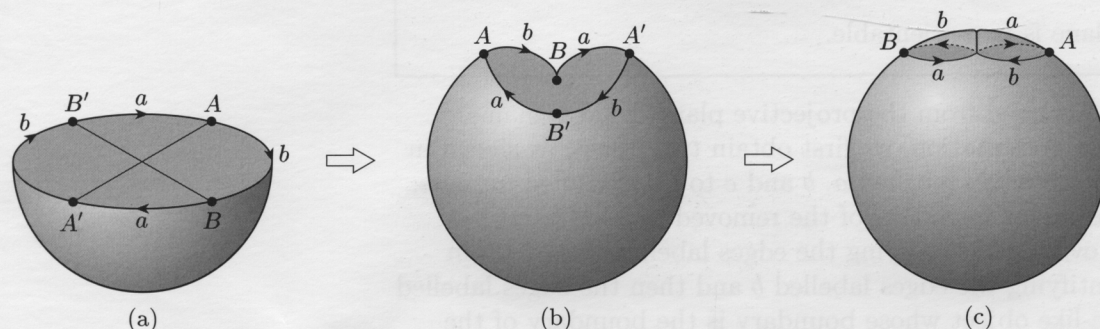


Figure 2.23

If we now remove a disc from the surface in Figure 2.23(c), as shown in Figure 2.24, Theorem 2.4 tells us that what we have left is a Möbius band. This new picture of a Möbius band is called a **cross-cap**, from its supposed resemblance to a bishop's mitre.

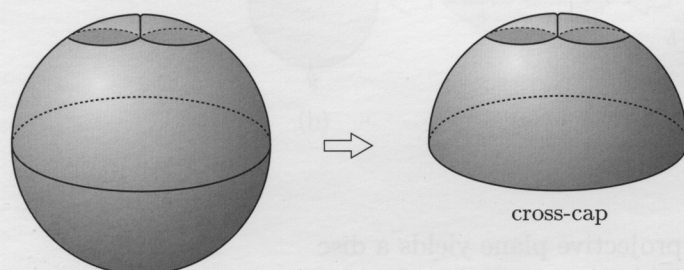


Figure 2.24

In general, we add a cross-cap to a surface by removing a disc and identifying opposite points on the boundary (or by identifying the boundary curve with that of a Möbius band). Adding one cross-cap to the sphere gives the projective plane; adding two cross-caps to the sphere gives the Klein bottle.

We stated in Section 1 that neither the projective plane nor the Klein bottle are surfaces in space. In both cases, whenever all identifications are made, the resulting representation is a surface that intersects itself and so cannot be embedded in \mathbb{R}^3 . In fact, a non-orientable surface can be represented as a surface in space if and only if it has a boundary.

Theorem 2.5

A non-orientable surface can be represented as a surface in space if and only if it has a boundary.

We omit the proof, which is beyond the scope of this unit.

We also have the following result.

Theorem 2.6

All orientable surfaces can be represented as surfaces in space.

We omit the proof, which is beyond the scope of this unit.

So the extra freedom that we gain by studying surfaces in general, and not just surfaces in space, is precisely the freedom to study non-orientable surfaces without boundary, of which the projective plane and the Klein bottle are the most important examples.

3 The Euler characteristic

After working through this section, you should be able to:

- ▶ explain the terms *subdivision* and *Euler characteristic*;
- ▶ draw a subdivision of a given surface and calculate its Euler characteristic;
- ▶ appreciate that the Euler characteristics of all subdivisions of a given surface are the same;
- ▶ state and use the Classification Theorem for surfaces.

In this section we introduce the third of the numbers we associate with a surface — the *Euler characteristic*. This is used in the Classification Theorem, which we state at the end of the section. To define the Euler characteristic, we need the idea of a *subdivision* of a surface, which we introduce by first considering the informal idea of a *net* on a surface.

3.1 Nets on surfaces

To define the Euler characteristic, we need the idea of covering a surface with a *net of polygons*. To explain what we mean let us look at some examples.

We start with a sphere and draw triangles upon it. One way to do this is to start with a tetrahedron, which (when considered as a polyhedral surface) is homeomorphic to the sphere. If we now ‘inflate’ the tetrahedron to obtain a sphere, the four triangular faces of the tetrahedron create a *net* of four triangles on the sphere, as Figure 3.1 illustrates.

Here the word ‘triangle’ means any three-sided region drawn on a surface. They usually have curved sides and are not flat.

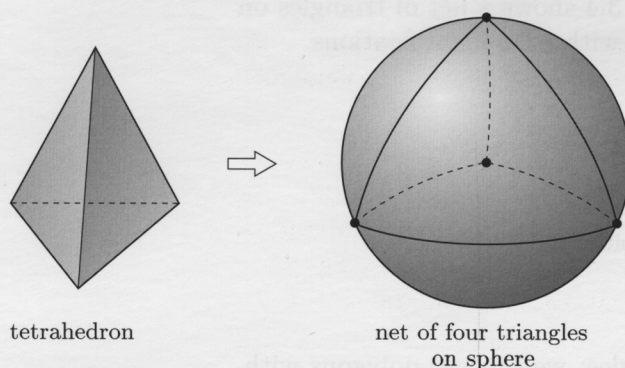


Figure 3.1

We can similarly inflate an octahedron, as shown in Figure 3.2.

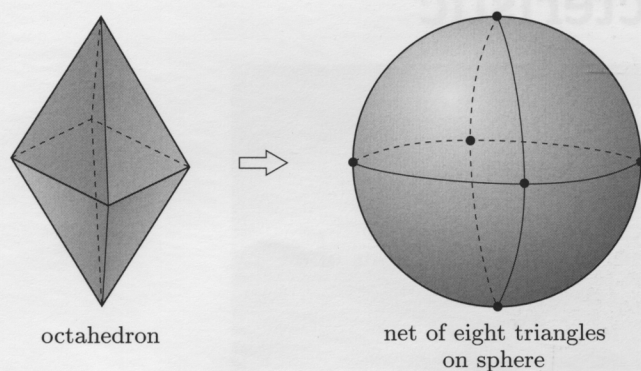


Figure 3.2

We can also draw nets of triangles on other surfaces, such as the torus, for which it is convenient to show the net on the representation of the torus as a rectangle with edge identifications, as in Figure 3.3.

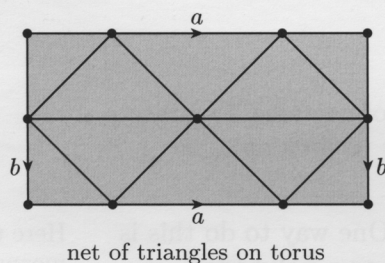


Figure 3.3

Remember that here the edge identifications a and b apply to the *whole* of the horizontal and vertical edges of the rectangle. For the sake of clarity, we assume that all parts of the edges of the defining polygon correspond to edges of the net.

When we draw triangles on a surface with boundary, unless stated otherwise, we require that the boundary is decomposed into edges that are edges of the triangles. For example, Figure 3.4 shows a net of triangles on a Möbius band, again shown as a rectangle with edge identifications.

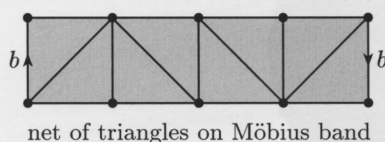


Figure 3.4

We need not restrict our attention to triangles: we can use polygons with any number of sides. For example, a cube can be inflated to give us a net of squares on a sphere, as Figure 3.5 shows.

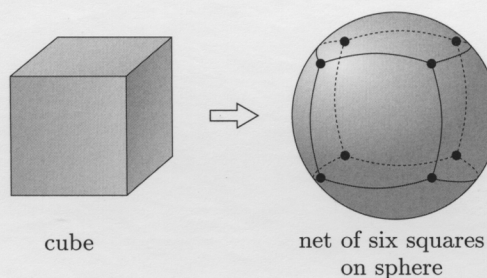


Figure 3.5

We can also use different types of polygon on the same surface. Figure 3.6 shows a net of triangles and quadrilaterals on a sphere, obtained by inflating a polyhedron called the *cuboctahedron*.

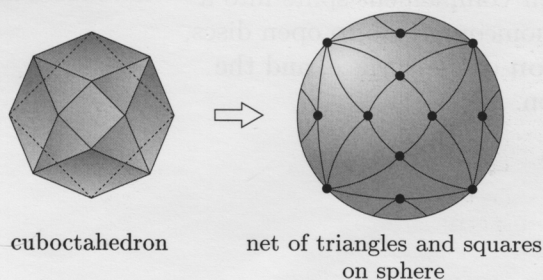


Figure 3.6

The surfaces we are considering include both surfaces in space and surfaces that cannot be represented in \mathbb{R}^3 , such as the Klein bottle and the projective plane. For example, Figure 3.7 shows a net of quadrilaterals on a Klein bottle, shown as a rectangle with edge identifications. Note that the faces that look like triangles are actually quadrilaterals, since each has four vertices and four edges.

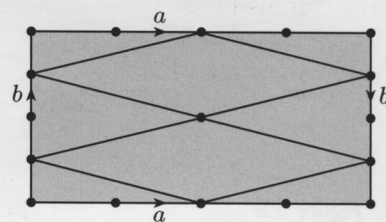


Figure 3.7

Problem 3.1

Draw an example of each of the following nets:

- a net of quadrilaterals on a torus;
- a net of triangles on a projective plane.

3.2 Subdivisions

In this subsection we formalize the idea of a net by introducing a useful concept, called a *subdivision* of a surface. This is a standard kind of net drawn on a surface, and is defined in terms of *vertices*, *edges* and *faces*. It leads to the idea of the *Euler characteristic* of the surface.

All surfaces obtained from polygons by identifying edges arise from a net (of sorts) consisting of a single polygonal face, together with the edges and vertices that remain after the identifications have been carried out. For example, a torus is obtained from a rectangle by identifying edges: the top and bottom edges become the same, the left and right edges become the same, the four vertices of the rectangle become the same, and there is just one face, as shown in Figure 3.8. This 'net' therefore has one vertex (A), two edges (a and b) and one face (F): the face F is what is left after all the vertices and edges have been removed.

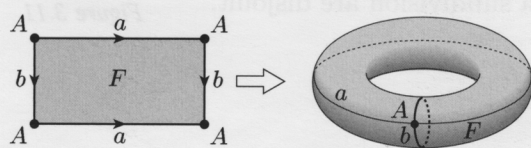


Figure 3.8

We now try to formalize the idea of a net of polygons on a surface so as to include both this case and the nets described in the previous subsection:

the result is called a *subdivision* of the surface. We start with the surface and mark a set of **vertices** (points) on it, joining them by **edges** (non-crossing curves) to produce a connected graph, as shown in Figure 3.9(a). If we now remove the vertices and edges of this graph from the surface as shown in Figure 3.9(b), then their complement splits into a finite number of pieces. If these pieces are all homeomorphic to open discs, then the vertices and edges define a **subdivision** of the surface, and the disc-like regions are the **faces** of the subdivision.

Connected graphs were introduced in Subsection 1.2.

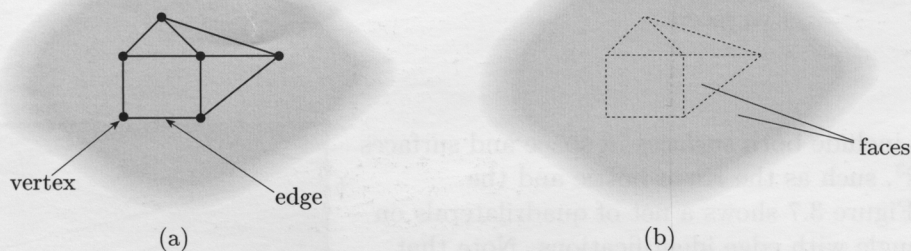


Figure 3.9

Before we can write down a formal definition of a subdivision of a surface, we must formalize the idea of an *edge*: an **edge** of a subdivision is the image of the open interval $(0, 1)$ under a one-one continuous map f of the closed interval $[0, 1]$, as Figure 3.10 illustrates. The points $f(0)$ and $f(1)$ are the **endpoints** of the edge.

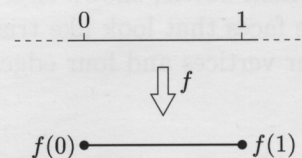


Figure 3.10

The endpoints are not part of the edge and may coincide.

Definition

A **subdivision** of a surface consists of a finite set of points, called **vertices**, and a finite set of **edges**, as defined above, such that:

- each vertex is an endpoint of at least one edge;
- each endpoint of an edge is a vertex;
- the vertices and edges form a connected graph;
- no two edges have any points in common;
- if the surface has a boundary, the boundary consists only of vertices and edges;
- the space obtained from the surface by removing the vertices and edges is a union of disjoint pieces, called **faces**, each of which is homeomorphic to an open disc.

An example of a subdivision is the 'cube subdivision' of a sphere, shown in Figure 3.11, which has 8 vertices, 12 edges and 6 faces.

It is natural to say that two adjacent faces 'meet along their common edge', and that edges 'meet at a vertex', but strictly speaking this is incorrect since the vertices, edges and faces of a subdivision are disjoint.

The vertices, edges and faces of a subdivision are disjoint.

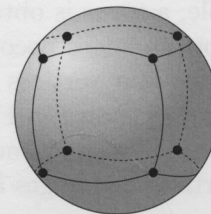


Figure 3.11

To allow us to use such natural phrases, we shall say that:

- a vertex *belongs to* an edge if every open set containing the vertex contains points of the edge, so the vertex is in the closure of the edge;
- a vertex *belongs to* a face if every open set containing the vertex contains points of the face, thus the vertex is in the closure of the face;
- an edge *belongs to* a face if every open set containing the edge contains points of the face, thus the edge is in the closure of the face.

These relationships are illustrated in Figure 3.12. They allow us to say that two edges or faces *meet* at a vertex if there is a vertex that belongs to both of them, and that two faces *meet* at an edge if there is an edge that belongs to both of them. They also allow us to say that a face *has* a number of edges and vertices, namely those edges and vertices that belong to it. We shall also say that the edges that meet at a vertex are the *incident edges* of that vertex.

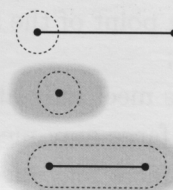


Figure 3.12

To illustrate this terminology, consider the subdivision of a closed disc shown in Figure 3.13, which has 5 vertices (A, B, C, D, E), 7 edges ($AB, AC, AD, AE, BC, BD, BE$) and 3 faces (1, 2, 3). Faces 1 and 2 meet at edges AD and BD and at vertices A, B and D ; faces 2 and 3 meet at edge AB and at vertices A and B ; faces 1 and 3 meet only at vertices A and B . Edges AC and BC meet at vertex C ; edges AC and AD meet at vertex A ; edges AD and BC do not meet; and so on. Face 1 has 4 edges (AC, AD, BC, BD) and 4 vertices (A, B, C, D); face 2 has 3 edges (AB, AD, BD) and 3 vertices (A, B, D); face 3 has 3 edges (AB, AE, BE) and 3 vertices (A, B, E).

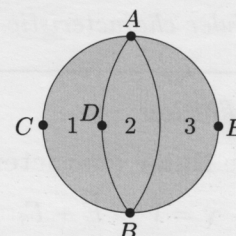


Figure 3.13

Problem 3.2

List the vertices, edges and faces of each of the subdivisions of a closed disc in Figure 3.14. Describe where face 1 meets the other faces in each case. How many edges and vertices does face 3 have in each case?

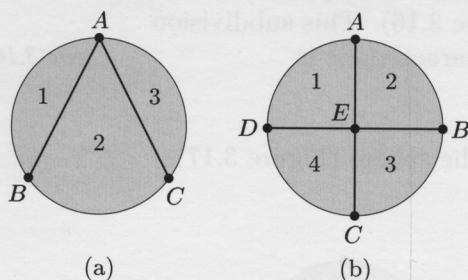


Figure 3.14

We noted earlier that the endpoints of an edge may coincide. This means that an edge can meet itself, as is illustrated by the circular edges in the subdivision of a closed disc shown in Figure 3.15. This figure also shows that a face may meet itself along an edge — face 2 meets itself along the edge AB (and of course at the vertices A and B) — and also a face may meet itself at a vertex or vertices but not along an edge — face 3 meets itself at vertices B and C only.

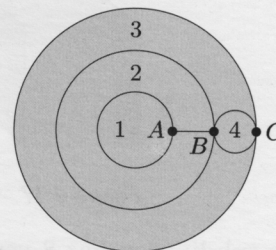


Figure 3.15

From the definition of a subdivision and the examples we have given, we may deduce the following properties of subdivisions of surfaces:

- each point of the surface must be a vertex, or lie on an edge, or lie in a face;
- faces meet only at vertices and edges;
- two faces may meet at one or more vertices;
- two faces may meet at one or more edges;
- edges meet only at vertices;
- an edge may belong to one or two faces;
- an edge may have one or two vertices;
- a vertex may belong to one or more faces;
- a vertex may belong to one or more edges.

3.3 The Euler characteristic

Subdivisions of surfaces lead to the third number used to classify surfaces, the *Euler characteristic*.

Definition

The **Euler characteristic** χ of a subdivision of a surface is

$$\chi = V - E + F,$$

where V , E and F denote (respectively) the numbers of vertices, edges and faces in the subdivision.

Remark

The formula $\chi = V - E + F$ is often referred to as **Euler's Formula**.

For example, the tetrahedron gives rise to a subdivision of the sphere where the vertices, edges and faces of the tetrahedron correspond to the vertices, edges and faces of the subdivision (Figure 3.16). This subdivision has 4 vertices, 6 edges and 4 faces, so its Euler characteristic is

$$\chi = V - E + F = 4 - 6 + 4 = 2.$$

Similarly, the cube gives rise to a subdivision of the sphere (Figure 3.17).

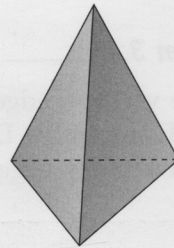


Figure 3.16

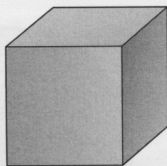


Figure 3.17

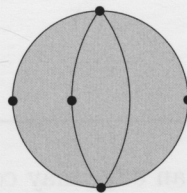


Figure 3.18

This subdivision has 8 vertices, 12 edges and 6 faces, so its Euler characteristic is

$$\chi = V - E + F = 8 - 12 + 6 = 2.$$

Figure 3.18 shows a subdivision of a closed disc, with 3 faces, 5 vertices and 7 edges. Its Euler characteristic is

$$\chi = 5 - 7 + 3 = 1.$$

We saw in the last subsection that the rectangle with edge identifications that define a torus (Figure 3.19) represents a subdivision of the torus, with 1 vertex, 2 edges and 1 face. Its Euler characteristic is

$$\chi = 1 - 2 + 1 = 0.$$

As the last example indicates, when we are dealing with a surface defined by a polygon with edge identifications, we need to be careful how we count the vertices, edges and faces. The following worked problem provides further illustration.

Worked problem 3.1

Find the Euler characteristic of the subdivision of the projective plane given by Figure 3.20.

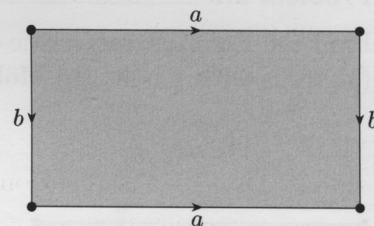


Figure 3.19

Further examples are given in the software.

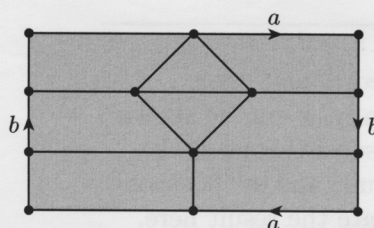


Figure 3.20

Solution

We label all the vertices, edges and faces, using the same labels for any that are identified, as in Figure 3.21.

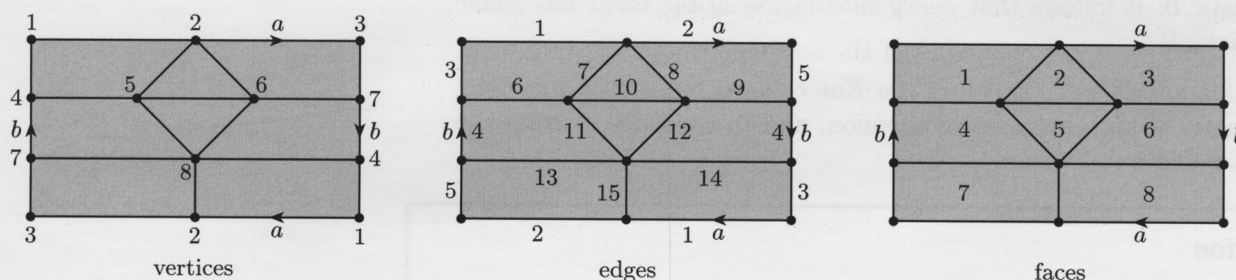


Figure 3.21

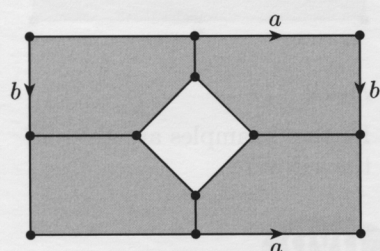
There are 8 vertices, 15 edges and 8 faces, so the Euler characteristic is

$$\chi = V - E + F = 8 - 15 + 8 = 1$$

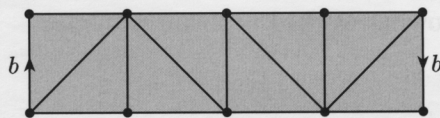
Problem 3.3

Find the Euler characteristic of each of the subdivisions in Figure 3.22:

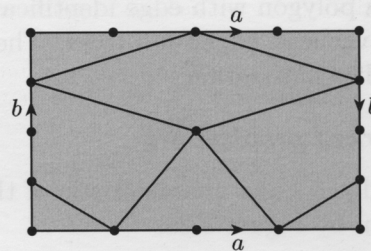
(a) torus with 1 hole, (b) Möbius band, (c) Klein bottle.



torus with 1 hole



Möbius band



Klein bottle

Figure 3.22

We have seen that the Euler characteristics of tetrahedral and cubic subdivisions of the sphere are equal to 2. More generally, it can be shown that the Euler characteristic is the same for *any two* subdivisions of the same surface. The proof of this general result is difficult, and we postpone a discussion of this issue until *Unit B3*. We simply state the result here.

Theorem 3.1

Any two subdivisions of the same surface have the same Euler characteristic.

For example, *every subdivision of the sphere has Euler characteristic 2*. Similarly, we have seen a subdivision of the torus with Euler characteristic 0: it follows that *every subdivision of the torus has Euler characteristic 0*.

Theorem 3.1 allows us to speak of the *Euler characteristic of a surface*, independently of the choice of subdivision, and to compute it using the most convenient subdivision.

Definition

The **Euler characteristic of a surface S** is the Euler characteristic of any subdivision of S . It is denoted by $\chi(S)$.

Our assumption that the surface is compact guarantees that the Euler characteristic is always finite.

χ is the Greek letter chi.

The earlier examples now enable us to conclude that the Euler characteristic of the sphere is 2, of the closed disc is 1, of the torus is 0, of the projective plane is 1, of the torus with 1 hole is -1 , of the Möbius band is 0 and of the Klein bottle is 0.

Historical note

A little history is instructive here, because it shows how difficult Theorem 3.1 really is. By 1900 the classification of compact surfaces was well understood, although proofs of the major theorems relied more on intuition than would be acceptable today. Attention switched to objects called '3-manifolds', topological spaces each point of which has a neighbourhood homeomorphic to the interior of the unit ball in \mathbb{R}^3 . The French mathematician Henri Poincaré — one of the leading mathematicians of his time — was among the first to study spaces of three or more dimensions. He proposed to study 3-manifolds by triangulating them — that is, by subdividing them into tetrahedra. Faced with the three-dimensional version of the problem of finding properties of a 3-manifold that are independent of any triangulation, he conjectured that given any two triangulations, there is a further triangulation which contains both. But not even he could prove it. His conjecture stood until 1948, and the corresponding result for \mathbb{R}^2 was proved only in 1963.

Perhaps the most eloquent indication of the difficulty of the proof of Theorem 3.1 was the remarkable discovery (by the American mathematician John Milnor in 1961) that the corresponding conjecture is false for spaces of dimension 6: this means that any proof in three dimensions must use some special property of three-dimensional space. This was one of a number of achievements for which Milnor received the Fields Medal in 1962, often regarded as the mathematical equivalent of the Nobel Prize.

Because of these difficulties, and even before Milnor's refutation of the general conjecture, a way had been found of avoiding it and working systematically in any number of dimensions. The method was pioneered in 1926 by another major mathematician, the Russian P. S. Alexandrov. His 600-page treatise *Combinatorial Topology* treats surfaces in Chapter III. You may find it consoling that he does not prove the invariance of the Euler characteristic in that chapter, but writes:

In contrast with the rest of the book, a great deal of the material in Chapter III depends on certain invariance theorems, which are not proved until Chapters V and X and are merely stated in this chapter ... This procedure seems to me legitimate, since the proofs of these invariance theorems given subsequently are independent of the results of Chapter III. Otherwise, the exposition of the topology of surfaces would have to be postponed until much later, or burdened with a proof of invariance for the special case of two dimensions.

In short, one of the creators of the modern theory thought it best to teach surfaces with less than complete rigour, rather than burden his readers with proofs that were different in spirit from the pure study of surfaces. (He gave a rigorous proof many pages later.) Indeed, no undergraduate textbook author presents a rigorous treatment that does not take the reader on a long journey, finding that an intuitive grasp of the topic is all that is needed at this level and that it is entirely reliable.

Problem 3.4

By considering the subdivisions in Figure 3.23, verify the Euler characteristic of each of the given surfaces.

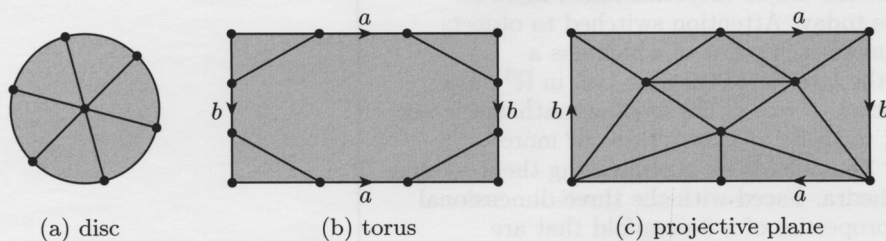


Figure 3.23

Problem 3.5

By drawing a suitable subdivision, find the Euler characteristic of a cylinder.

We conclude this subsection by remarking that any homeomorphism of a surface S onto a surface S' maps a subdivision of S onto a subdivision of S' , sending the vertices of S to vertices of S' , the edges of S to edges of S' , and the faces of S to faces of S' , in a one-to-one fashion. It follows that the value of $V - E + F$ remains unchanged, and therefore that the Euler characteristic of a subdivision is preserved under homeomorphisms. Thus the Euler characteristic of a surface is a topological invariant.

3.4 Some general results

We next establish some general results about Euler characteristics. We start with a theorem that tells us what happens to the Euler characteristic of a surface when we remove an open disc.

Theorem 3.2 Disc Removal Theorem

The Euler characteristic of a surface with an open disc removed is one less than the Euler characteristic of the original surface.

Outline proof Suppose that we are given a surface S containing an open disc. Any subdivision of S with the disc removed gives rise to a subdivision of S with the disc reinstated, as shown in Figure 3.24, and these two subdivisions have the same number V of vertices and the same number E of edges.

The full proof is rather complicated, so we provide just an outline.

Recall that, if a surface has a boundary, then the boundary must consist of vertices and edges of the subdivision.

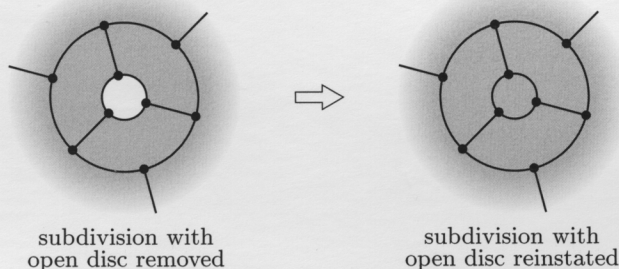


Figure 3.24

When we reinstate the disc, we increase F , the number of faces, by 1, while V and E are unchanged. So we increase the Euler characteristic $\chi = V - E + F$ by 1. It follows that the Euler characteristic of the surface with the disc removed is one less than the Euler characteristic of the original surface. ■

Surfaces with holes

Using this result, we can obtain the Euler characteristic of a surface with any number of holes, by successively inserting the holes one at a time. For example, since a closed disc has Euler characteristic 1, it follows that a closed disc with 1 hole has Euler characteristic 0, a disc with 2 holes has Euler characteristic -1 , and so on.

Problem 3.6

Use the above approach to determine the Euler characteristic of a torus with k holes.

Theorem 3.3

Inserting k holes into a surface reduces its Euler characteristic by k .

n -fold toruses

We can use a similar technique to find the Euler characteristic of a 2-fold torus. If we cut the surface into two, as shown in Figure 3.25, and separate the pieces, we obtain two copies of a 1-fold torus with 1 hole, each with Euler characteristic -1 .

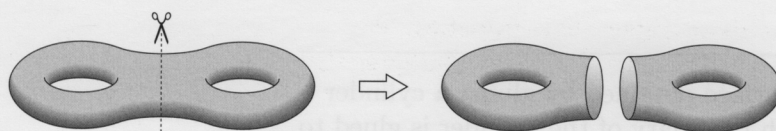


Figure 3.25

Consider any subdivision of the left-hand torus with 1 hole, and give the right-hand torus with 1 hole its mirror image subdivision. In this way, the vertices and edges on the boundary of each hole can be identified when we glue the two pieces back together. Note that the number of vertices on each boundary must be the same as the number of edges on that boundary.

Thus the vertices and edges on the boundaries contribute a value 0 to the Euler characteristic $V - E + F$ of each piece, and also to the Euler characteristic of the 2-fold torus obtained by gluing the two pieces back together. Thus, the Euler characteristic of the 2-fold torus is simply the sum of the Euler characteristics of the two pieces: $(-1) + (-1) = -2$.

We can carry out a similar procedure for a 3-fold torus. If we cut it into two pieces, as shown in Figure 3.26, we obtain a 2-fold torus with 1 hole and a 1-fold torus with 1 hole. The Euler characteristic of a 2-fold torus with 1 hole is $(-2) - 1 = -3$, by Theorem 3.2, and the Euler characteristic of a 1-fold torus with 1 hole is -1 , so the Euler characteristic of a 3-fold torus is $(-3) + (-1) = -4$.

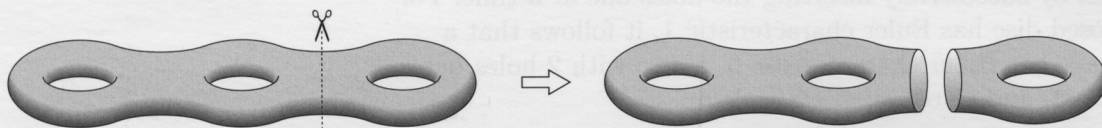


Figure 3.26

Continuing in this way, we obtain the Euler characteristic of an n -fold torus.

Theorem 3.4

The Euler characteristic of an n -fold torus is $2 - 2n$.

We can also deduce the following theorem from the above discussion.

Theorem 3.5

The Euler characteristic of a surface obtained by gluing together two surfaces with boundary along a given boundary component of each is equal to the sum of the Euler characteristics of the original surfaces.

Problem 3.7

Find the Euler characteristic of the surface obtained by gluing a cylinder to a sphere with two holes, where each boundary of the cylinder is glued to the boundary of a separate hole.

3.5 The Classification Theorem

In this subsection we state the Classification Theorem for surfaces, which classifies a surface in terms of its boundary number β , its orientability number ω and its Euler characteristic χ , each of which is a topological invariant — it is preserved under homeomorphisms.

Let us remind ourselves of these three numbers.

- A surface may or may not have a boundary, and, if it does, then the boundary has finitely many disjoint pieces. The number of these boundary components is the boundary number β of the surface. Subsection 1.2.
- Some surfaces are orientable and some are not. The orientability number ω of a surface is 0 if the surface is orientable and 1 if it is non-orientable. Subsection 2.2.
- Surfaces can be subdivided, and for a given surface the Euler characteristic $\chi = V - E + F$ is independent of the subdivision used to compute it, and is therefore a property of the surface itself. Subsection 3.3.

These three numbers β , ω and χ are collectively called the **characteristic numbers** of the surface.

Now, all three characteristic numbers are necessary for classifying surfaces since, for example:

- a sphere and a torus both have $\beta = 0$ and $\omega = 0$;
- a torus and a Klein bottle both have $\beta = 0$ and $\chi = 0$;
- a torus and a cylinder both have $\omega = 0$ and $\chi = 0$.

It also turns out that these three numbers are sufficient to classify surfaces, and so we have the central result of this block.

Theorem 3.6 Classification Theorem for surfaces

Two surfaces are homeomorphic if and only if they have the same values for the characteristic numbers β , ω and χ .

Remarks

- (i) We already know that the characteristic numbers are topological invariants, that is, any two homeomorphic surfaces have the same values for the characteristic numbers. Thus it is solely the converse, namely if two surfaces have the same values for the characteristic numbers then they are homeomorphic, that we have to prove.
- (ii) It follows from the Classification Theorem that if we restrict our attention to *orientable* surfaces ($\omega = 0$), then the numbers β and χ characterize the surface completely — and similarly if we restrict our attention to non-orientable surfaces ($\omega = 1$).
- (iii) Our assumption that all surfaces are compact guarantees that the characteristic numbers are finite.
- (iv) The Classification Theorem allows us to refer to *the* sphere, *the* torus, *the* Klein bottle, etc., rather than a sphere, a torus and a Klein bottle, etc.

The proof of the converse, and hence of the Classification Theorem, is given in *Unit B3*.

Problem 3.8

Which surfaces are characterized by the following sets of characteristic numbers?

- (a) $\beta = 0$, $\omega = 0$, $\chi = 2$.
- (b) $\beta = 1$, $\omega = 1$, $\chi = 0$.
- (c) $\beta = 1$, $\omega = 0$, $\chi = -1$.
- (d) $\beta = 0$, $\omega = 0$, $\chi = -2$.
- (e) $\beta = 0$, $\omega = 1$, $\chi = 1$.

Problem 3.9

Find the characteristic numbers of the surface obtained by gluing two Möbius bands along their boundaries. Deduce from the Classification Theorem that this surface is homeomorphic to the Klein bottle.

Worked problem 3.2

Find the characteristic numbers of the closed disc with two straps, shown in Figure 3.27.

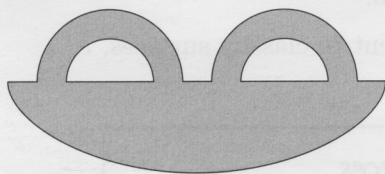


Figure 3.27

Solution

The surface has three boundary components, so $\beta = 3$.

The surface contains no Möbius band, and is therefore orientable, so $\omega = 0$.

To find the Euler characteristic, we choose a subdivision. The one shown in Figure 3.28 has 8 vertices, 12 edges and 3 faces, so $\chi = 8 - 12 + 3 = -1$.

You could also argue that, as the surface is a closed disc with two holes, its Euler characteristic is $1 - 2 = -1$ (by Theorem 3.3).

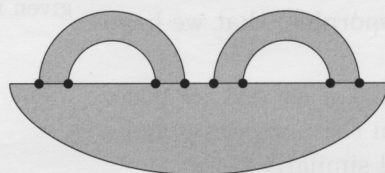


Figure 3.28

Problem 3.10

Consider the closed disc with two straps added in other ways, sometimes with a twist, as shown in Figure 3.29. Compute the characteristic numbers of each surface, and then use the Classification Theorem to determine which of these surfaces are homeomorphic.

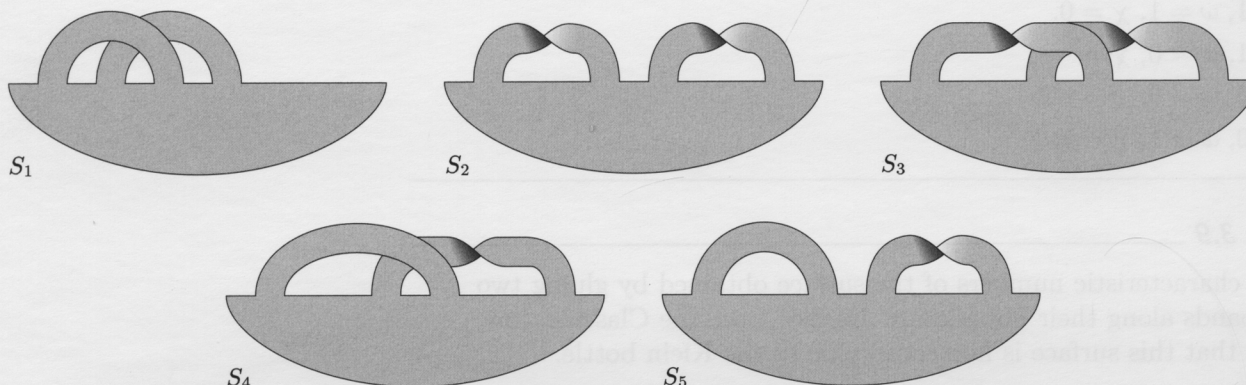


Figure 3.29

Problem 3.11

For each of the following paper surfaces, mark a subdivision and hence determine the characteristic numbers of the surface.

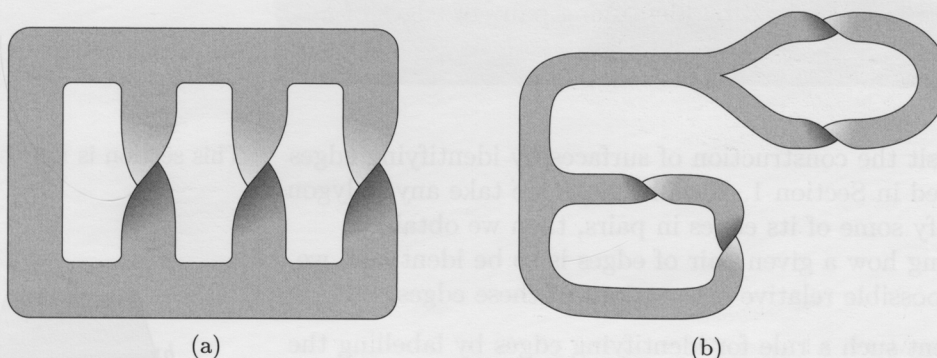


Figure 3.30

For future reference, we conclude this section with a table of some common surfaces with their characteristic numbers.

surface	β	ω	χ
sphere	0	0	2
n -fold torus	0	0	$2 - 2n$
Klein bottle	0	1	0
projective plane	0	1	1
closed disc	1	0	1
cylinder	2	0	0
Möbius band	1	1	0

In addition, adding a hole to any surface increases its boundary number β by 1, keeps its orientability number ω the same, and decreases its Euler characteristic χ by 1.

4 Edge identifications

After working through this section, you should be able to:

- explain the term *identification topology*;
- explain why an object obtained by identifying pairs of edges of a polygon is a surface.

In this section, we revisit the construction of surfaces by identifying edges of polygons, as described in Section 1. Recall that, if we take any polygon in the plane and identify some of its edges in pairs, then we obtain a surface. When specifying how a given pair of edges is to be identified, we choose one of the two possible relative orientations of these edges.

In diagrams we represent such a rule for identifying edges by labelling the edges with letters and arrowheads, as in Figure 4.1: a pair of edges is identified if each carries the same letter, and the direction of identification is indicated by the arrowheads. We refer to an edge labelled with an arrowhead as a *directed edge*. Any assignment of letters and arrowheads to edges in which a particular letter appears on exactly two directed edges defines a rule for identifying edges of a polygon.

To *prove* that we do obtain a surface when we identify edges of a polygon in this way, we need to show how to make the resulting set into a topological space. We use a topology called the *identification topology*. We explore this topology and its consequences in this section.

This section is not assessed.

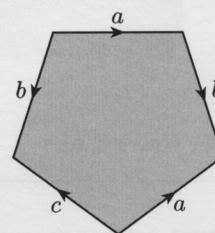


Figure 4.1

4.1 The identification topology

Our aim is to show that the object that we produce when we identify some or all the edges of a polygon is a surface. Therefore, by the definition of a surface given in Subsection 1.5, we must show how it can be given the structure of a topological space, and that this space is Hausdorff. Furthermore, we must show that every point has a disc-like or half-disc-like neighbourhood. We shall address the issue of showing that the object is a topological space in this subsection. We shall consider neighbourhoods and show that the space is Hausdorff in Subsection 4.2.

Our method for showing that the object is a topological space will be to construct a bijection between it and a known topological space. This will then allow us to define open sets for the object and hence show that it is a topological space.

The obvious known topological space is the polygon with edge identifications used to construct the object, since any polygon can be considered to have the subspace topology inherited from the Euclidean topology on \mathbb{R}^3 . We also have a map from the polygon to the object, the map that takes points of the polygon to points of the object.

Unfortunately, this map is not one-one, as it maps points on the edges to be identified to the same point on the object. However, we can turn this into a one-one map by creating what is known as the *identification space* of the polygon. Then, if we can show that this identification space is a topological space, we shall have essentially achieved our objective.

Its compactness and path-connectedness are guaranteed by its method of construction.

The map is onto.

To define the identification space, we first need to define what we mean by an *identification class*. Consider the identification of opposite edges of a rectangle to form a cylinder, as illustrated in Figure 4.2.

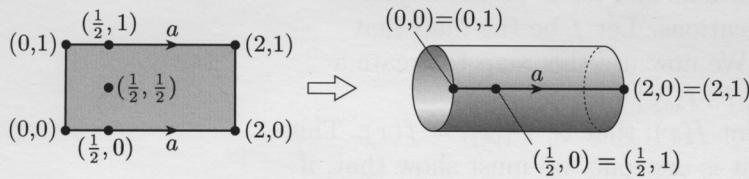


Figure 4.2

For each point x of the rectangle, we define the *identification class* of x to be the set of all points of the rectangle that are mapped to the same point of the cylinder as x is; for example, the class of $(\frac{1}{2}, 0)$ consists of the pair of points $(\frac{1}{2}, 0)$ and $(\frac{1}{2}, 1)$, while the class of $(\frac{1}{2}, \frac{1}{2})$ consists of just the point $(\frac{1}{2}, \frac{1}{2})$. In general, we have the following definition.

Definition

Let $f : X \rightarrow Y$ be a function and let $x \in X$. The **identification class** of x , denoted by $[x]$, is the set

$$[x] = \{z \in X : f(z) = f(x)\}.$$

Worked problem 4.1

Consider the identification of opposite edges of the rectangle $[0, 3] \times [0, 2]$ to form a torus, as illustrated in Figure 4.3. Write down the identification class $[x]$ of each of the following points:

- (a) $x = (2, 1)$; (b) $x = (2, 0)$; (c) $x = (3, 1)$; (d) $x = (3, 0)$.

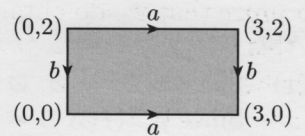


Figure 4.3

Solution

- (a) $[x] = \{(2, 1)\}$;
 (b) $[x] = \{(2, 0), (2, 2)\}$;
 (c) $[x] = \{(3, 1), (0, 1)\}$;
 (d) $[x] = \{(3, 0), (0, 0), (0, 2), (3, 2)\}$.

Problem 4.1

Consider the identification of opposite edges of the rectangle $[0, 3] \times [0, 2]$ to form a Klein bottle, as illustrated in Figure 4.4. Write down the identification class $[x]$ of each of the following points:

- (a) $x = (2, 1)$; (b) $x = (2, 0)$; (c) $x = (3, \frac{1}{2})$; (d) $x = (3, 0)$.

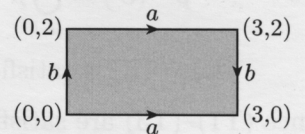


Figure 4.4

This now allows us to define the identification space of a polygon with edge identifications. In general the definition is as follows.

Definition

Given a topological space (X, \mathcal{T}) , a set Y , and a continuous function $f : X \rightarrow Y$, the **identification space** $I(X)$ of X under f is the set of all identification classes $[x]$, for $x \in X$.

Let X be a polygon with edge identifications and let Y be the object obtained by making those edge identifications. Let f be the map that creates Y from X , so that $f(X) = Y$. We now use this map to create a bijection F from $I(X)$ to Y , by letting $F: I(X) \rightarrow Y$ map the identification class $[x]$ to the image point $f(x)$, that is, $F([x]) = f(x)$. This function is clearly onto. To show that it is one-one, we must show that, if $F([a]) = F([b])$, then $[a] = [b]$. This follows since, if $F([a]) = F([b])$, then $f(a) = f(b)$, and thus $[a] = [b]$. Thus F is a bijection, as required.

The next step is to define a topology on $I(X)$. We first note that there is an obvious map p from X onto $I(X)$, sending each element x to its identification class $[x]$. We now define a subset U of $I(X)$ to be open if $p^{-1}(U)$ is an open subset of X . This then creates the desired topology on $I(X)$, as the following result demonstrates.

Theorem 4.1

Let (X, \mathcal{T}_X) be a topological space, let Y be a set, let $f: X \rightarrow Y$ be a function, let $I(X)$ be the identification space of X under f and let $[x]$ be the identification class of $x \in X$. Let $p: X \rightarrow I(X)$ be the function defined by $p(x) = [x]$. Then $\mathcal{T}_f = \{U \subseteq I(X) : p^{-1}(U) \in \mathcal{T}_X\}$ is a topology on $I(X)$.

Proof We check that \mathcal{T}_f satisfies conditions (T1)–(T3) for a topology.

(T1) Since $p(\emptyset) = \emptyset$, $\emptyset = p^{-1}(\emptyset) \in \mathcal{T}_X$, and so $\emptyset \in \mathcal{T}_f$.

Since $p^{-1}(I(X)) = X \in \mathcal{T}_X$, $I(X) \in \mathcal{T}_f$.

Thus (T1) is satisfied.

(T2) Let $U_1, U_2 \in \mathcal{T}_f$ and let $U = U_1 \cap U_2$.

Then $p^{-1}(U_1) \in \mathcal{T}_X$ and $p^{-1}(U_2) \in \mathcal{T}_X$, and so

$$p^{-1}(U) = p^{-1}(U_1 \cap U_2) = p^{-1}(U_1) \cap p^{-1}(U_2) \in \mathcal{T}_X.$$

Thus (T2) is satisfied.

(T3) Let $\{U_i : i \in I\}$ be a family of sets in \mathcal{T}_f , and let $U = \bigcup_{i \in I} U_i$.

Then $p^{-1}(U_i) \in \mathcal{T}_X$ for each $i \in I$, and so

$$p^{-1}(U) = \bigcup_{i \in I} p^{-1}(U_i) \in \mathcal{T}_X.$$

Thus (T3) is satisfied.

Since (T1)–(T3) are satisfied, \mathcal{T}_f is a topology on $I(X)$. ■

Thus $(I(X), \mathcal{T}_f)$ is a topological space. We give the topology \mathcal{T}_f a special name.

Definition

The topology \mathcal{T}_f defined in Theorem 4.1 is the **identification topology** for X under f .

We now use F to create a topology on $Y = f(X)$. We define a subset U of Y to be open if $F^{-1}(U)$ is an open subset of $I(X)$. Recall that $F([x]) = f(x)$.

Problem 4.2

Show that the family \mathcal{T}_Y of open sets defined in this way comprises a topology on Y .

So we have a bijection $F: I(X) \rightarrow Y$ between the topological spaces $(I(X), \mathcal{T}_f)$ and (Y, \mathcal{T}_Y) . Furthermore, from the definitions of open sets in \mathcal{T}_f and \mathcal{T}_Y , we can deduce immediately that, for each $U \in \mathcal{T}_Y$, $F^{-1}(U) \in \mathcal{T}_f$ and, for each $V \in \mathcal{T}_f$, $F(V) \in \mathcal{T}_Y$. Hence F is a homeomorphism between $(I(X), \mathcal{T}_f)$ and (Y, \mathcal{T}_Y) .

We have thus shown that the object created by identifying some or all of the edges of a polygon is a topological space homeomorphic to the identification space of the polygon. We now need to show that this topological space is Hausdorff and that every point has a disc-like or half-disc-like neighbourhood, to demonstrate that the space is indeed a surface.

4.2 Neighbourhoods

We know that a polygon X is a surface and so each point x in X has a disc-like or half-disc-like neighbourhood. We shall show that a map f that identifies edges of a polygon to create an object Y automatically creates corresponding disc-like or half-disc-like neighbourhoods of each point $y = f(x)$ of Y .

If x is in the interior of X , there is no difficulty: the point x has a disc-like neighbourhood U which is mapped by f to a disc-like neighbourhood $f(U)$ of $y = f(x)$ in Y .

If x lies on the boundary of X , there are three cases to consider:

- x lies on an edge that is identified with another edge, but not at a vertex of that edge;
- x lies at a vertex of an edge that is identified with another edge;
- x lies on a part of the boundary that is not identified with another part.

In each case, x has a half-disc-like neighbourhood in X . Some examples will help us determine which happens to these half-disc-like neighbourhoods after identification takes place.

Torus

If x lies on an edge, then each of the two points in $[x]$ has a half-disc-like neighbourhood. When we identify edges, these neighbourhoods fit together to form disc-like neighbourhoods on the torus, as Figure 4.5 shows.

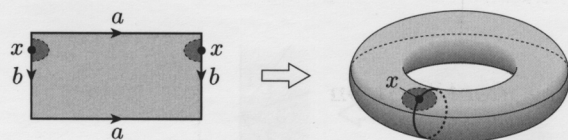


Figure 4.5

If x is a vertex, then each of the four points in $[x]$ has a half-disc-like neighbourhood (in the shape of a quarter-disc). When we identify edges, these neighbourhoods fit together to form a disc-like neighbourhood on the torus, as Figure 4.6 shows.

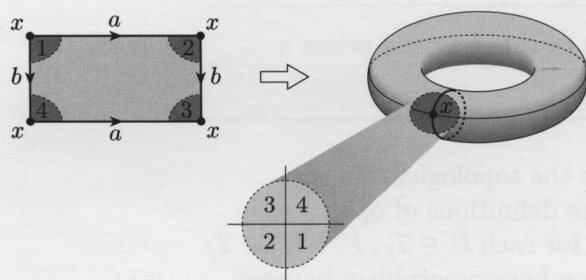


Figure 4.6

Klein bottle

If x lies on an edge, then each of the two points in $[x]$ has a half-disc-like neighbourhood (see Figure 4.7). When we identify edges, these neighbourhoods fit together to form disc-like neighbourhoods in the Klein bottle.

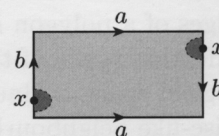


Figure 4.7

If x is a vertex, then each of the four points in $[x]$ has a half-disc-like neighbourhood (as the shape of a quarter-disc — see Figure 4.8). When we identify edges, these neighbourhoods fit together to form a disc-like neighbourhood in the Klein bottle.

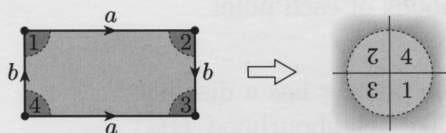


Figure 4.8

Torus with 1 hole

In our last example, we consider a pentagon with two pairs of edges identified. As we saw in Subsection 1.3, identification of the edges produces a torus with a hole. In this case there are five vertex-neighbourhoods to fit together, as shown in Figure 4.9. Identification of edges a and b shows that neighbourhoods 2, 3 and 4 fit together as shown. Identification of the edges b shows that neighbourhood 5 fits alongside neighbourhood 2, and identification of edges a shows that neighbourhood 1 fits alongside neighbourhood 4. However, neighbourhoods 1 and 5 do not fit together, and so we do not obtain a disc-like neighbourhood — rather we obtain a half-disc-like neighbourhood and the resulting vertex lies on the boundary of the hole.

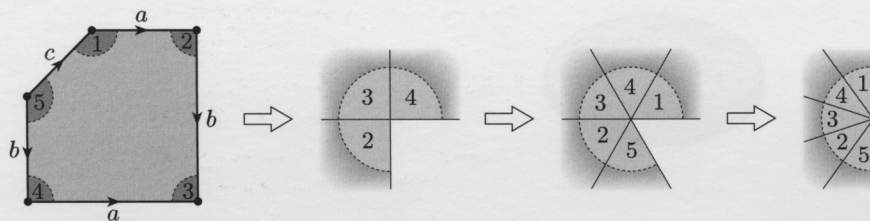


Figure 4.9

Finally consider a point x that lies on the edge c of the pentagon (see Figure 4.10). This has a half-disc-like neighbourhood. After identification of edges, this point lies on the boundary of the hole, and maintains its half-disc-like neighbourhood.

From these examples we can deduce that:

- if a point lies on an edge that is identified with another edge, but not at a vertex of that edge, then after identification that point has a disc-like neighbourhood;
- if a point lies on a part of the boundary that is not identified with another part, then after identification that point has a half-disc-like neighbourhood.

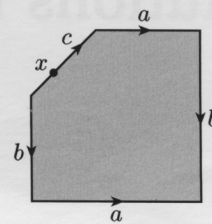


Figure 4.10

However, if a point is vertex of an edge that is identified with another edge, we have seen that after identification that point can have a disc-like or half-disc-like neighbourhood. We need to be sure that these are the only possibilities.

Consider the vertices x_1, x_2, \dots, x_k of a polygon that are identified to the same point x by the edge identifications. In the polygon, each vertex x_i has a half-disc-like neighbourhood that can be represented as in Figure 4.11. We shall call such a piece a *wedge*.



Figure 4.11

Some of the edges that come into each vertex x_i are to be identified. Let us see what can happen. First, notice that the edge identifications that involve the vertices x_i cannot have wedges that fit together to form two or more separate pieces: this is because, in such a case, the edge identifications would not identify all the vertices x_i to a single point x . Also, because whole edges and not just vertices are identified, we cannot produce a situation such as that illustrated in Figure 4.12. Thus, the edge identifications that involve the vertices x_i are such that all the wedges fit together to form a single neighbourhood. Either the point x is in the interior of the object so formed, as for the torus and Klein bottle, giving rise to a disc-like neighbourhood formed from the wedges, or the point x lies on the boundary, as for the torus with 1 hole. Thus:

- if a point lies at a vertex of an edge that is identified with another edge, then after identification that point has a disc-like or half-disc-like neighbourhood.

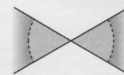
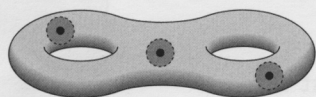


Figure 4.12

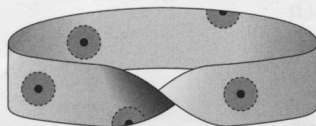
We have therefore demonstrated that every point of polygon with edge identifications has, after identification, a disc-like or half-disc-like neighbourhood. Furthermore, given any two points of the object obtained after identification, by choosing disc-like or half-disc-like neighbourhoods of sufficiently small radius, we can find neighbourhoods that do not intersect — so the object is Hausdorff. Combining these results with those of Subsection 4.1 tells us that an object formed by identifying edges of a polygon is a surface, as defined in Subsection 1.5. Furthermore, we know that such a surface is homeomorphic to the identification space of the polygon with edge identifications.

Solutions to problems

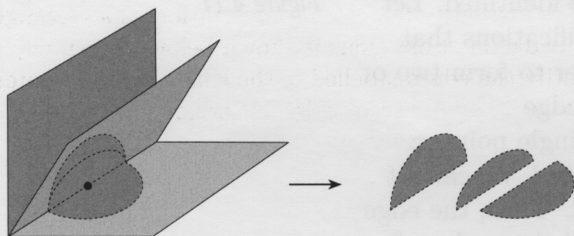
1.1



1.2

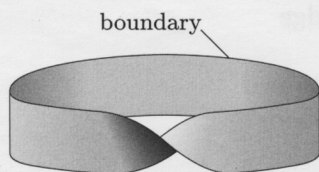


1.3 No point on the spine of a 3-leafed book has a disc-like or half-disc-like neighbourhood. If we remove the spine, then the neighbourhood of any point on the spine breaks into three pieces, as shown. This cannot be the case if the neighbourhood is homeomorphic to a disc or half-disc.



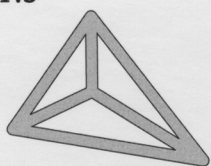
Since no point on the spine has a disc-like or half-disc-like neighbourhood, the 3-leafed book is not a surface.

1.4 (a) The boundary consists of the single edge.

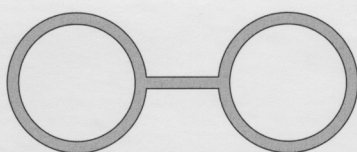


(b) Since each point has a disc-like neighbourhood, there is no boundary, i.e. the boundary is the empty set.

1.5



3-fold torus



2-fold torus

1.6



4-fold torus with 5 holes

The boundary consists of the boundaries of the 5 holes.

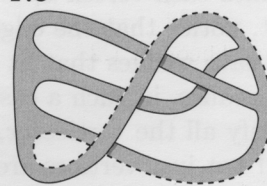
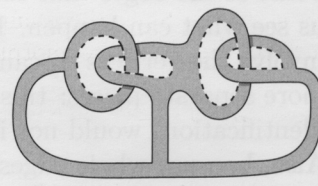


5-fold torus with 4 holes

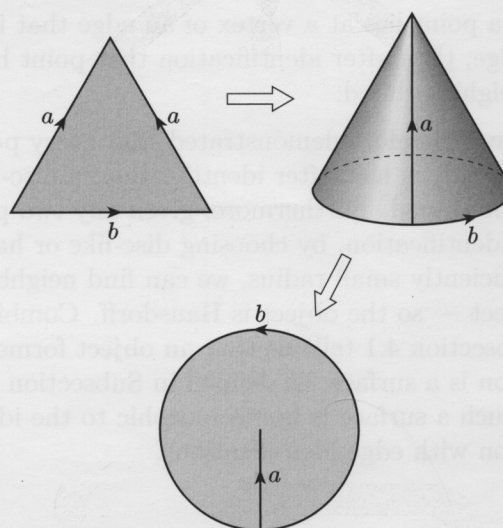
The boundary consists of the boundaries of the 4 holes.

1.7 (a) $\beta = 1$; (b) $\beta = 4$; (c) $\beta = 1$.

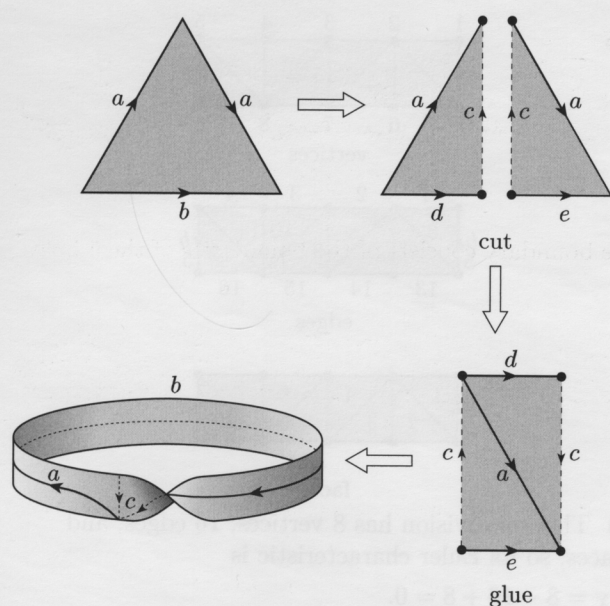
1.8

(a) $\beta = 2$ (b) $\beta = 5$

1.9 (a) Identifying the edges we obtain a cone which can then be flattened out. The result is a disc.



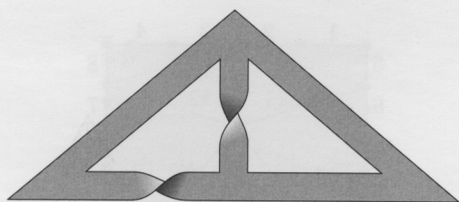
(b) The three vertices of the triangle are all identified to a single point which lies on the boundary of the resulting surface; the two edges that are identified become a single closed curve through this point. The result is a *Möbius band*. One way of seeing this is to bisect the triangle, glue the two parts together by identifying the edges labelled *a*, to create a rectangle, then repair the bisection (which needs a twist).



1.10 For the Möbius band, $\beta = 1$; for the 2-fold torus, $\beta = 0$. Since these numbers are different, the surfaces are not homeomorphic, by Theorem 1.2.

2.1 (a) $\beta = 1$. (b) $\beta = 2$.

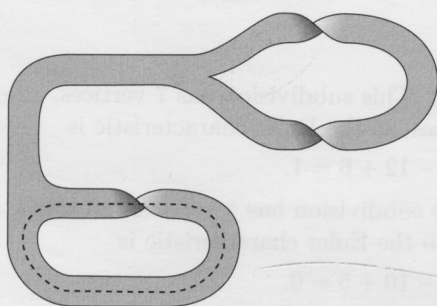
2.2 One way to insert a half-twist to obtain a surface with boundary number $\beta = 1$ is the following.



2.3 (a) $\omega = 0$.

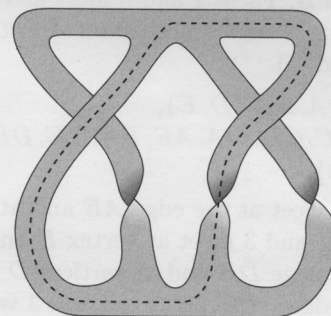
(b) A band with three half-twists is homeomorphic to a Möbius band, and so, by Theorem 2.1, $\omega = 1$.

2.4 (a) This surface contains a Möbius band, indicated by the closed curve, so is non-orientable ($\omega = 1$).



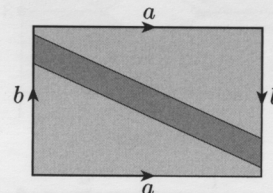
(b) This surface contains no Möbius bands, so is orientable ($\omega = 0$).

(c) This surface contains several Möbius bands, one of which is indicated by the closed curve, so is non-orientable ($\omega = 1$).

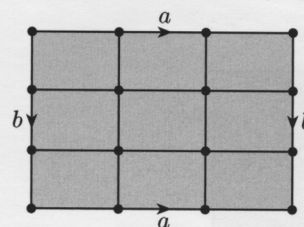


2.5 $(3, -1, 2)$, $(-9, 3, -6)$ and $(15, -5, 10)$ correspond to the projective point $[3, -1, 2]$. $(4, -2, 6)$ and $(-18, 9, -27)$ correspond to the projective point $[4, -2, 6] = [2, -1, 3]$. $(12, -6, -9)$ corresponds to the projective point $[12, -6, -9] = [4, -2, -3]$.

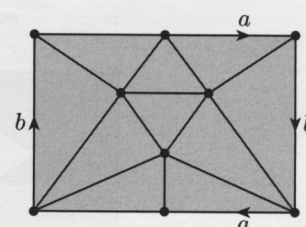
2.6 The Klein bottle's representation as a rectangle with edge identifications is shown below. When we identify the edges labelled b , the shaded strip becomes a Möbius band, and thus by Theorem 2.2 the Klein bottle is non-orientable.



3.1 Examples of suitable nets are the following. In each case we draw the net on the surface represented as a rectangle with edge identifications.



(a)



(b)

3.2 (a) vertices $\{A, B, C\}$,
edges $\{AB \text{ (twice)}, AC \text{ (twice)}, BC\}$;
faces $\{1, 2, 3\}$.

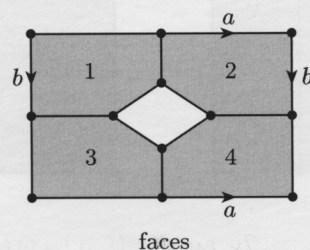
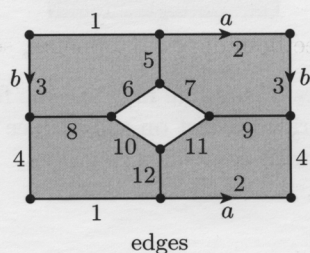
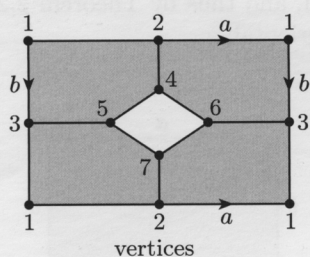
Faces 1 and 2 meet at the inner edge AB and at vertices A and B . Faces 1 and 3 meet at vertex A only. Face 3 has 2 edges (both from A to C) and 2 vertices (A and C).

(b) vertices $\{A, B, C, D, E\}$,
edges $\{AB, BC, CD, DA, AE, BE, CE, DE\}$,
faces $\{1, 2, 3, 4\}$.

Faces 1 and 2 meet at the edge AE and at vertices A and E . Faces 1 and 3 meet at vertex E only. Faces 1 and 4 meet at edge DE and at vertices D and E . Face 3 has 3 edges (BC, BE, CE) and 3 vertices (B, C and E).

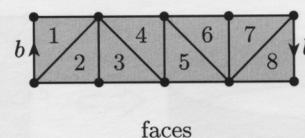
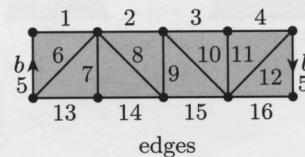
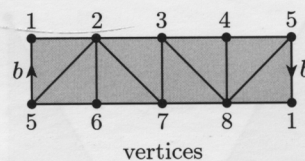
3.3 (a) This subdivision has 7 vertices, 12 edges, and 4 faces, so its Euler characteristic is

$$\chi = 7 - 12 + 4 = -1.$$



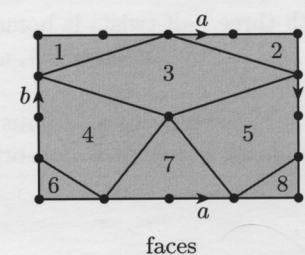
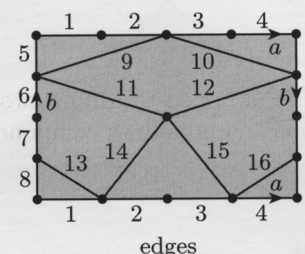
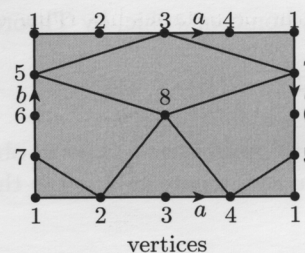
(b) This subdivision has 8 vertices, 16 edges, and 8 faces, so its Euler characteristic is

$$\chi = 8 - 16 + 8 = 0.$$



(c) This subdivision has 8 vertices, 16 edges, and 8 faces, so its Euler characteristic is

$$\chi = 8 - 16 + 8 = 0.$$



3.4 (a) This subdivision has 7 vertices, 12 edges, and 6 faces, so the Euler characteristic is

$$\chi = 7 - 12 + 6 = 1.$$

(b) This subdivision has 5 vertices, 10 edges, and 5 faces, so the Euler characteristic is

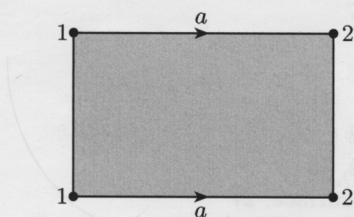
$$\chi = 5 - 10 + 5 = 0.$$

(c) This subdivision has 6 vertices, 15 edges, and 10 faces, so the Euler characteristic is

$$\chi = 6 - 15 + 10 = 1.$$

3.5 A simple subdivision is shown below. It has 2 vertices, 3 edges and 1 face, so the Euler characteristic is

$$\chi = 2 - 3 + 1 = 0.$$



3.6 Since a torus has Euler characteristic 0, it follows from Theorem 3.2 that a torus with 1 hole has Euler characteristic -1 , a torus with 2 holes has Euler characteristic -2 , and so on. In general, a torus with k holes has Euler characteristic $-k$.

3.7 A cylinder has Euler characteristic 0 and a sphere with two holes has Euler characteristic $2 - 2 = 0$ (by Theorem 3.3). Therefore, by Theorem 3.5, the surface obtained by gluing these two surfaces together along the boundaries has Euler characteristic $0 + 0 = 0$.

3.8 (a) Sphere.

(b) Möbius band.

(c) Torus with 1 hole.

(d) 2-fold torus.

(e) Projective plane.

3.9 The surface obtained by gluing two Möbius bands together has no boundary components ($\beta = 0$), and is non-orientable ($\omega = 1$). Also, by Theorem 3.5, its Euler characteristic is the sum of the Euler characteristics of each separate surface, so $\chi = 0 + 0 = 0$. The Klein bottle has these values for its characteristic numbers, so (by the Classification Theorem) the glued surface is homeomorphic to the Klein bottle.

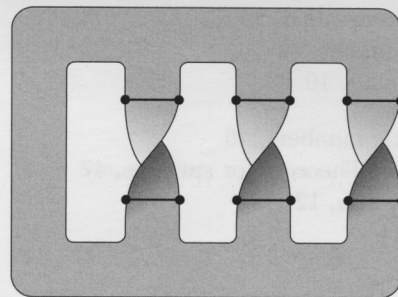
3.10 Each surface has a subdivision with 8 vertices, 12 edges and 3 faces, as indicated in Worked problem 3.2. The Euler characteristic of each surface is therefore -1 , as before.

The characteristic numbers β and ω are as follows.

surface	S_1	S_2	S_3	S_4	S_5
β	1	1	2	1	2
ω	0	1	1	1	1

It follows from the Classification Theorem that the surfaces S_2 and S_4 are homeomorphic, as are the surfaces S_3 and S_5 ; no other two surfaces are homeomorphic.

3.11 (a) We have drawn a subdivision with 12 vertices.

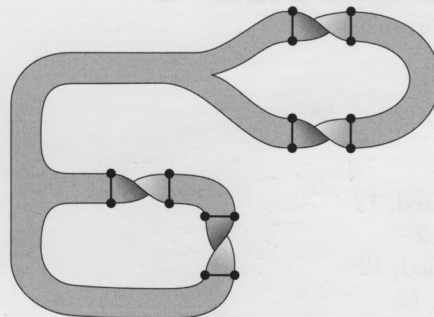


There are 18 edges and 4 faces, so the Euler characteristic is $\chi = 12 - 18 + 4 = -2$.

There is one boundary component, so $\beta = 1$.

The surface is non-orientable (see Worked problem 2.1(b)), so $\omega = 1$.

(b) We have drawn a subdivision with 16 vertices.



There are 24 edges and 7 faces, so the Euler characteristic is $\chi = 16 - 24 + 7 = -1$.

There are three boundary components, so $\beta = 3$.

The surface is orientable (see Solution 2.4(b)), so $\omega = 0$.

4.1 (a) $[x] = \{(2, 1)\}$;

(b) $[x] = \{(2, 0), (2, 2)\}$;

(c) $[x] = \{(3, \frac{1}{2}), (0, \frac{3}{2})\}$;

(d) $[x] = \{(3, 0), (0, 0), (0, 2), (3, 2)\}$.

4.2 We check (T1)–(T3).

(T1) Since $F(\emptyset) = \emptyset$, $\emptyset = F^{-1}(\emptyset) \in \mathcal{T}_f$, and so $\emptyset \in \mathcal{T}_Y$. Since $F^{-1}(Y) = I(X) \in \mathcal{T}_f$, $Y \in \mathcal{T}_Y$. Thus (T1) is satisfied.

(T2) Let $U_1, U_2 \in \mathcal{T}_Y$ and let $U = U_1 \cap U_2$. Then $F^{-1}(U_1) \in \mathcal{T}_f$ and $F^{-1}(U_2) \in \mathcal{T}_f$, and so

$$F^{-1}(U) = F^{-1}(U_1 \cap U_2) = F^{-1}(U_1) \cap F^{-1}(U_2) \in \mathcal{T}_f.$$

Thus (T2) is satisfied.

(T3) Let $\{U_i : i \in I\}$ be a family of sets in \mathcal{T}_Y and let $U = \bigcup_{i \in I} U_i$. Then $F^{-1}(U_i) \in \mathcal{T}_f$ for each $i \in I$, and so $F^{-1}(U) = \bigcup_{i \in I} F^{-1}(U_i) \in \mathcal{T}_f$.

Thus (T3) is satisfied.

Since (T1)–(T3) are satisfied, \mathcal{T}_Y is a topology on Y .

Index

- boundary, 10
- boundary component, 14
- boundary number, 14
- boundary point, 10

- characteristic numbers, 46
- Classification Theorem for surfaces, 47
- connected graph, 12
- cross-cap, 34
- cube, 7
- cylinder, 8, 15

- Disc Removal Theorem, 44
- disc-like neighbourhood, 7, 9

- edge, 12, 38
 - of graph, 12
 - of subdivision, 38
- edge identifications, 19, 31, 50
- Euler characteristic, 40
- Euler characteristic of surface, 42
- Euler's Formula, 40

- face, 38
 - of subdivision, 38

- graph, 12
 - connected, 12
 - edge, 12
 - thickened, 12
 - vertex, 12

- half-disc-like neighbourhood, 8, 9
- half-twist, 24
- Hausdorff space, 22
- homeomorphism, 4

- identification, 15
- identification class, 51
- identification space, 51
- identification topology, 52

- identifying edges, 16, 17, 23, 50
- incident edges, 39

- Klein bottle, 17

- Möbius band, 8, 16

- n -fold torus, 18
- net of polygons, 35
- non-orientable surface, 27

- open half-disc, 8
- orientability number, 28
- orientable surface, 27

- path-connected surface, 22
- projective plane, 17, 30
- projective point, 31

- sphere, 7, 18
- subdivision, 37, 38
- surface, 9, 10, 22–24
 - compact, 23
 - path-connected, 22
 - with boundary, 10
 - with twists, 24
 - without boundary, 10
- surface in space, 10, 23, 27

- thickened graph, 12
- topological invariant, 21
- torus, 7, 13, 16–18
 - n -fold, 11, 18
 - with holes, 13, 17

- vertex, 12, 37
 - of graph, 12
 - of subdivision, 37

- whole-twist, 24